arXiv:math/0110124v4 [math.AG] 1 Jan 2003

A p-adic local monodromy theorem

Kiran S. Kedlaya
Department of Mathematics
University of California, Berkeley
Berkeley, CA 94720
kedlaya@math.berkeley.edu

December 31, 2002

Abstract

We produce a canonical filtration for locally free sheaves on an open p-adic annulus equipped with a Frobenius structure. Using this filtration, we deduce a conjecture of Crew on p-adic differential equations, analogous to Grothendieck's local monodromy theorem (also a consequence of results of André and of Mebkhout). Namely, given a finite locally free sheaf on an open p-adic annulus with a connection and a compatible Frobenius structure, the module admits a basis over a finite cover of the annulus on which the connection acts via a nilpotent matrix.

Contents

1 Introduction **2**
1.1 Crew's conjecture on p-adic local monodromy 2
1.2 Frobenius filtrations and Crew's conjecture 4
1.3 Applications 5
1.4 Structure of the paper 6
1.5 An example: the Bessel isocrystal 7

2 A few rings **8**
2.1 Notations and conventions 9
2.2 Valued fields 9
2.3 The "classical" case $K = k((t))$ 11
2.4 More on Bézout rings 14
2.5 σ-modules and (σ, ∇)-modules 16

3 A few more rings **18**

3.1 Cohen rings 18
3.2 Overconvergent rings 20
3.3 Analytic rings: generalizing the Robba ring 24
3.4 Some σ-equations 30
3.5 Factorizations over analytic rings 34
3.6 The Bézout property for analytic rings 41

4 The special Newton polygon 45
4.1 Properties of eigenvectors 45
4.2 Existence of eigenvectors 48
4.3 Raising the Newton polygon 51
4.4 Construction of the special Newton polygon 57

5 The generic Newton polygon 59
5.1 Properties of eigenvectors 60
5.2 The Dieudonné-Manin classification 61
5.3 Slope filtrations 64
5.4 Comparison of the Newton polygons 67

6 From a slope filtration to quasi-unipotence 69
6.1 Approximation of matrices 69
6.2 Some matrix factorizations 72
6.3 Descending the special slope filtration 74
6.4 The connection to the unit-root case 79
6.5 Logarithmic form of Crew's conjecture 82

1 Introduction

1.1 Crew's conjecture on p-adic local monodromy

The role of p-adic differential equations in algebraic geometry was first pursued systematically by Dwork; the modern manifestation of this role comes via the theory of isocrystals and F-isocrystals, which over a field of characteristic $p > 0$ attempt to play the part of local systems for the classical topology on complex varieties and lisse sheaves for the l-adic topology when $l \neq p$. In order to get a usable theory, however, an additional "overconvergence" condition must be imposed, which has no analogue in either the complex or l-adic cases. For example, the cohomology of the affine line is infinite dimensional if computed using convergent isocrystals, but has the expected dimension if computed using overconvergent isocrystals. This phenomenon was generalized by Monsky and Washnitzer [MW] into a cohomology theory for smooth affine varieties, and then generalized further by Berthelot to the theory of rigid cohomology, which has good behavior for arbitrary varieties (see for example [Be1]).

Unfortunately, the use of overconvergent isocrystals to date has been hampered by a gap in the local theory of these objects; for example, it obstructed the proof of finite dimensionality of Berthelot's rigid cohomology with arbitrary coefficients (the case of constant coefficients was treated by Berthelot in [Be2]). This gap can be described as a p-adic analogue of Grothendieck's local monodromy theorem for l-adic cohomology.

The best conceivable analogue of Grothendieck's theorem would be that an F-isocrystal becomes a successive extension of trivial isocrystals after a finite étale base extension. Unfortunately, this assertion is not correct; for example, it fails for the pushforward of the constant isocrystal on a family of ordinary elliptic curves degenerating to a supersingular elliptic curve (and for the Bessel isocrystal described in Section 1.5 over the affine line).

The correct analogue of the local monodromy theorem was formulated conjecturally by Crew [Cr2, Section 10.1], and reformulated in a purely local form by Tsuzuki [T2, Theorem 5.2.1]; we now introduce some terminology and notation needed to describe it. (These definitions are reiterated more precisely in Chapter 2.) Let k be a field of characteristic $p > 0$, and $\mathcal{O}$ a finite totally ramified extension of a Cohen ring $C(k)$. The Robba ring $\Gamma_{\mathrm{an,con}}$ is defined as the set of Laurent series over $\mathcal{O}[\frac{1}{p}]$ which converge on some open annulus with outer radius 1; its subring Γ_{con} consists of series which take integral values on some open annulus with outer radius 1, and is a discrete valuation ring. (See Chapter 3 to find out where the notations come from.) We say a ring endomorphism $\sigma : \Gamma_{\mathrm{an,con}} \to \Gamma_{\mathrm{an,con}}$ is a *Frobenius* for $\Gamma_{\mathrm{an,con}}$ if it is a composition power of a map preserving Γ_{con} and reducing modulo a uniformizer of Γ_{con} to the p-th power map. For example, one can choose $t \in \Gamma_{\mathrm{con}}$ whose reduction is a uniformizer in the ring of Laurent series over k, then set $t^\sigma = t^q$. Note that one cannot hope to define a Frobenius on the ring of analytic functions on any fixed open annulus with outer radius 1, because for η close to 1, functions on the annulus of inner radius 1 pull back under σ to functions on the annulus of inner radius $\eta^{1/p}$. Instead, one must work over an "infinitely thin" annulus of radius 1.

Given a ring R in which $p \neq 0$ and an endomorphism $\sigma : R \to R$, we define a *σ-module* M as a finite locally free module equipped with an R-linear map $F : M \otimes_{R,\sigma} R \to M$ that becomes an isomorphism over $R[\frac{1}{p}]$; the tensor product notation indicates that R is viewed as an R-module via σ. For the rings considered in this paper, a finite locally free module is automatically free; see Proposition 2.5. Then F can be viewed as an additive, σ-linear map $F : M \to M$ that acts on any basis of M by a matrix invertible over $R[\frac{1}{p}]$.

We define a *(σ, ∇)-module* as a σ-module plus a connection $\nabla : M \to M \otimes \Omega^1_{R/\mathcal{O}}$ (that is, an additive map satisfying the Leibniz rule $\nabla(c\mathbf{v}) = c\nabla(\mathbf{v}) + \mathbf{v} \otimes dc$) that makes the following diagram commute:

$$\begin{tikzcd}
M \arrow[r, "\nabla"] \arrow[d, "F"] & M \otimes \Omega^1_{R/\mathcal{O}} \arrow[d, "F \otimes d\sigma"] \\
M \arrow[r, "\nabla"] & M \otimes \Omega^1_{R/\mathcal{O}}
\end{tikzcd}$$

We say a (σ, ∇)-module over $\Gamma_{\mathrm{an,con}}$ is *quasi-unipotent* if, after tensoring $\Gamma_{\mathrm{an,con}}$ over Γ_{con} with a finite extension of Γ_{con}, the module admits a filtration by (σ, ∇)-submodules such that each successive quotient admits a basis of elements in the kernel of ∇. (If k is perfect, one may also insist that the extension of Γ_{con} be residually separable.) In these notations, Crew's conjecture is resolved by the following theorem, which we will prove in a more precise form as Theorem 6.12.

Theorem 1.1 (Local monodromy theorem). *Let σ be any Frobenius for the Robba ring $\Gamma_{\mathrm{an,con}}$. Then every (σ, ∇)-module over $\Gamma_{\mathrm{an,con}}$ is quasi-unipotent.*

Briefly put, a p-adic differential equation on an annulus with a Frobenius structure has quasi-unipotent monodromy. It is worth noting (though not needed in this paper) that for a given ∇, whether there exists a compatible F does not depend on the choice of the Frobenius map σ. This follows from the existence of change of Frobenius functors [T2, Theorem 3.4.10].

The purpose of this paper is to establish some structural results on modules over the Robba ring yielding a proof of Theorem 1.1. Note that Theorem 1.1 itself has been established independently by André [A2] and by Mebkhout [M]. However, as we describe in the next section, the methods in this paper are essentially orthogonal to the methods of those authors. In fact, the different approaches provide different auxiliary information, various pieces of which may be of relevance in other contexts.

1.2 Frobenius filtrations and Crew's conjecture

Before outlining our approach to Crew's conjecture, we describe by way of contrast the common features of the work of André and Mebkhout. Both authors build upon the results of a series of papers by Christol and Mebkhout [CM1], [CM2], [CM3], [CM4] concerning properties of modules with connection over the Robba ring. Most notably, in [CM4] is produced a canonical filtration (the "weight filtration"), defined whether or not the connection admits a Frobenius structure. André and Mebkhout show (in two different ways) that when a Frobenius structure is present, the graded pieces of this filtration can be shown to be quasi-unipotent.

The strategy in this paper is in a sense completely orthogonal to the aforementioned approach. (For a more detailed comparison between the various approaches to Crew's conjecture, see the November 2001 Seminaire Bourbaki talk of Colmez [Co].) Instead of isolating the connection data, we isolate the Frobenius structure and prove a structure theorem for σ-modules over the Robba ring. This can be accomplished by a "big rings" argument, where one first proves that σ-modules can be trivialized over a large auxiliary ring, and then "descends" the construction back to the Robba ring. (Isolating Frobenius is in a sense natural from the point of view of crystalline cohomology; for example, this is the approach of Katz in [Ka].)

The model for our strategy of trivializing σ-modules over an auxiliary ring is the Dieudonné-Manin classification of σ-modules over a complete discrete valuation ring R of mixed characteristic $(0, p)$ with algebraically closed residue field. (This classification is a semilinear analogue of the diagonalization of matrices over an algebraically closed field, except that here there is no failure of semisimplicity.) We give a quick statement here, deferring the precise formulation to Section 5.2. For $\lambda \in \mathcal{O}[\frac{1}{p}]$ and d a positive integer, let $M_{\lambda,d}$ denote the σ-module of rank d over $R[\frac{1}{p}]$ on which F acts by a basis $\mathbf{v}_1, \dots, \mathbf{v}_d$ as follows:

$$\begin{aligned} F\mathbf{v}_1 &= \mathbf{v}_2 \\ &\vdots \\ F\mathbf{v}_{d-1} &= \mathbf{v}_d \\ F\mathbf{v}_d &= \lambda \mathbf{v}_1. \end{aligned}$$

Define the *slope* of $M_{\lambda,d}$ to be $v_p(\lambda)/d$. Then the Dieudonné-Manin classification states (in part) that over $R[\frac{1}{p}]$, every σ-module is isomorphic to a direct sum $\oplus_j M_{\lambda_j,d_j}$, and the slopes that occur do not depend on the decomposition.

If R is a discrete valuation ring of mixed characteristic $(0,p)$, we may define the slopes of a σ-module over $R[\frac{1}{p}]$ as the slopes in a Dieudonné-Manin decomposition over the maximal unramified extension of the completion of R. However, this definition cannot be used immediately over $\Gamma_{\mathrm{an,con}}$, because that ring is not a discrete valuation ring. Instead, we must first reduce to considering modules over Γ_{con}. Our main theorem makes it possible to do so. Again, we give a quick formulation here and prove a more precise result later as Theorem 6.10. (Note: the filtration in this theorem is similar to what Tsuzuki [T2] calls a "slope filtration for Frobenius structures".)

Theorem 1.2. *Let M be a σ-module over $\Gamma_{\mathrm{an,con}}$. Then there is a canonical filtration $0 = M_0 \subset M_1 \subset \cdots \subset M_l = M$ of M by saturated σ-submodules such that:*

(a) each quotient M_i/M_{i-1} is isomorphic over $\Gamma_{\mathrm{an,con}}$ to a σ-module N_i defined over $\Gamma_{\mathrm{con}}[\frac{1}{p}]$;

(b) the slopes of N_i are all equal to some rational number s_i;

(c) $s_1 < \cdots < s_l$.

The relevance of this theorem to Crew's conjecture is that (σ, ∇)-modules over $\Gamma_{\mathrm{con}}[\frac{1}{p}]$ with a single slope can be shown to be quasi-unipotent using a result of Tsuzuki [T1]. The essential case is that of a unit-root (σ, ∇)-module over Γ_{con}, in which all slopes are 0. Tsuzuki showed that such modules become isomorphic to a direct sum of trivial (σ, ∇)-modules after a finite base extension, and even gave precise information about what extension is needed. This makes it possible to deduce the local monodromy theorem from Theorem 1.2.

1.3 Applications

We now describe some consequences of the results of this paper, starting with some applications via Theorem 1.1. One set of consequences occurs in the study of Berthelot's rigid cohomology (a sort of "grand unified theory" of p-adic cohomologies). For example, Theorem 1.1 can be used to establish finite dimensionality of rigid cohomology with coefficients in an overconvergent F-isocrystal; see [Cr2] for the case of a curve and [Ke7] for the general case. It can also be used to generalize the results of Deligne's "Weil II" to overconvergent F-isocrystals; this is carried out in [Ke8], building on work of Crew [Cr1], [Cr2]. In addition, it can be used to treat certain types of "descent", such as Tsuzuki's full faithfulness conjecture [T3], which asserts that convergent morphisms between overconvergent F-isocrystals are themselves overconvergent; see [Ke6].

Another application of Theorem 1.1 has been found by Berger [Bg], who has exposed a close relationship between F-isocrystals and p-adic Galois representations. In particular, he shows that Fontaine's "conjecture de monodromie p-adique" for p-adic Galois representations (that every de Rham representation is potentially semistable) follows from Theorem 1.1.

Further applications of Theorem 1.2 exist that do not directly pass through Theorem 1.1. For example, André (work in progress) has formulated a q-analogue of Crew's conjecture, in which the single differential equation is replaced by a formal deformation. He has established this analogue using Theorem 6.10 plus a q-analogue of Tsuzuki's unit-root theorem (Proposition 6.11), and has deduced a finiteness theorem for rigid cohomology of q-F-isocrystals. (It should also be possible to obtain these results using a q-analogue of the Christol-Mebkhout theorem, and indeed André and di Vizio have made progress in this direction; however, at the time of this writing, some technical details had not yet been worked out.)

We also plan to establish, in a subsequent paper, a conjecture of Shiho [Sh, Conjecture 3.1.8], on extending overconvergent F-isocrystals to log-F-isocrystals after a generically étale base change. This result appears to require a more sophisticated analogue of Theorem 6.10, in which the "one-dimensional" Robba ring is replaced by a "higher-dimensional" analogue. (One might suspect that this conjecture should follow from Theorem 1.1 and some clever geometric arguments, but the situation appears to be more subtle.) Berthelot (private communication) has suggested that a suitable result in this direction may help in constructing Grothendieck's six operations in the category of arithmetic $\mathcal{D}$-modules, which would provide a p-adic analogue of the constructible sheaves in étale cohomology.

1.4 Structure of the paper

We now outline the strategy of the proof of Theorem 1.2, and in the process describe the structure of the paper. We note in passing that some of the material appears in the author's doctoral dissertation [Ke1], written under Johan de Jong, and/or in a sequence of unpublished preprints [Ke2], [Ke3], [Ke4], [Ke5]. However, the present document avoids any logical dependence on unpublished results.

In Chapter 2, we recall some of the basic rings of the theory of p-adic differential equations; they include the Robba ring, its integral subring and the completion of the latter (denoted the "Amice ring" in some sources). In Chapter 3, we construct some less familiar rings by augmenting the classical constructions. These augmentations are inspired by (and in some cases identical to) the auxiliary rings used by de Jong [dJ] in his extension to equal characteristic of Tate's theorem [Ta] on p-divisible groups over mixed characteristic discrete valuation rings. (They also resemble the "big rings" in Fontaine's theory of p-adic Galois representations.) In particular, a key augmentation, denoted $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$, is a sort of "maximal unramified extension" of the Robba ring, and a great effort is devoted to showing that it shares the Bézout property with the Robba ring; that is, every finitely generated ideal in $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ is principal. (This chapter is somewhat technical; we suggest that the reader skip it on first reading, and refer back to it as needed.)

With these augmented rings in hand, in Chapter 4 we show that every σ-module over the Robba ring can be equipped with a canonical filtration over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$; this amounts to an "overconvergent" analogue of the Dieudonné-Manin classification. From this filtration we read off a sequence of slopes, which in case we started with a quasi-unipotent (σ, ∇)-module agree with the slopes of Frobenius on a nilpotent basis; the Newton polygon with these slopes is called the *special Newton polygon*.

By contrast, in Chapter 5, we associate to a (σ, ∇)-module over Γ_{con} the Frobenius slopes produced by the Dieudonné-Manin classification. The Newton polygon with these slopes is called the *generic Newton polygon*. Following [dJ], we construct some canonical filtrations associated with the generic Newton polygon. This chapter is logically independent of Chapter 4 except at its conclusion, when the two notions of Newton polygon are compared. In particular, we show that the special Newton polygon lies above the generic Newton polygon with the same endpoint, and obtain additional structural consequences in case the Newton polygons coincide.

Finally, in Chapter 6, we take the "generic" and "special" filtrations, both defined over large auxiliary rings, and descend them back to the Robba ring itself. The basic strategy here is to separate positive and negative powers of the series parameter, using the auxiliary filtrations to guide the process. Starting with a σ-module over the Robba ring, the process yields a σ-module over Γ_{con} whose generic and special Newton polygons coincide. The structural consequences mentioned above yield Theorem 1.2; by applying Tsuzuki's theorem on unit-root (σ, ∇)-modules (Proposition 6.11), we deduce a precise form of Theorem 1.1.

1.5 An example: the Bessel isocrystal

To clarify the remarks of the previous section, we include a classical example to illustrate the different structures we have described, especially the generic and special Newton polygons. Our example is the Bessel isocrystal, first studied by Dwork [Dw]; our description is a summary of the discussion of Tsuzuki [T2, Example 6.2.6] (but see also André [A1]).

Let p be an odd prime, and put $\mathcal{O} = \mathbb{Z}_p[\pi]$, where π is a $(p-1)$-st root of $-p$. Choose $\eta < 1$, and let R be the ring of Laurent series over $\mathcal{O}$ convergent for $|t| > \eta$ in the variable. Let σ be the Frobenius lift on $\mathcal{O}$ such that $t^\sigma = t^p$. Then for suitable η, there exists a (σ, ∇)-module M of rank two over R admitting a basis $\mathbf{v}_1, \mathbf{v}_2$ such that

$$\begin{aligned} F\mathbf{v}_1 &= A_{11}\mathbf{v}_1 + A_{12}\mathbf{v}_2 \\ F\mathbf{v}_2 &= A_{21}\mathbf{v}_1 + A_{22}\mathbf{v}_2 \\ \nabla\mathbf{v}_1 &= t^{-2}\pi^2\mathbf{v}_2 \otimes dt \\ \nabla\mathbf{v}_2 &= t^{-1}\mathbf{v}_1 \otimes dt. \end{aligned}$$

Moreover, the matrix A satisfies

$$\det(A) = p \qquad \text{and} \qquad A \equiv \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \pmod{p}.$$

It follows that the two generic Newton slopes are nonnegative (because the entries of A are integral), their sum is 1 (by the determinant equation), and the smaller of the two is zero (by the congruence). Thus the generic Newton slopes are 0 and 1.

On the other hand, M becomes unipotent after adjoining a square root of t to the Robba ring. More precisely, if $y = (t/4)^{1/2}$, define

$$f_\pm = 1 + \sum_{n=1}^{\infty} (\pm 1)^n \frac{(1 \times 3 \times \cdots \times (2n-1))^2}{(8\pi)^n n!} y^n$$

and set

$$\mathbf{w}_\pm = f_\pm \mathbf{e}_1 + \left(y \frac{df_\pm}{dy} + \left(\frac{1}{2} \mp \pi y^{-1} \right) f_\pm \right) \mathbf{e}_2.$$

Then

$$\nabla \mathbf{w}_\pm = \left(\frac{-1}{2} \pm \pi y^{-1} \right) \mathbf{w}_\pm \otimes \frac{dy}{y}.$$

Using the compatibility between the Frobenius and connection structures, we deduce that

$$F\mathbf{w}_\pm = \alpha_\pm y^{-(p-1)/2} \exp(\pm \pi (y^{-1} - y^{-\sigma}))$$

for some $\alpha_+, \alpha_- \in \mathcal{O}[\frac{1}{p}]$ with $\alpha_+ \alpha_- = 2^{1-p} p$. By the invariance of Frobenius under the automorphism $y \to -y$ of $\Gamma_{\mathrm{an,con}}[y]$, we deduce that α_+ and α_- have the same valuation.

It follows from this discussion that M is unipotent over $\Gamma_{\mathrm{an,con}}[y]$, and the two slopes of the special Newton polygon are equal, necessarily to $1/2$ since their sum is 1. In particular, the special Newton polygon lies above the generic Newton polygon and has the same endpoint, but the two polygons are not equal in this case.

Acknowledgments

The author was supported by a Clay Mathematics Institute Liftoffs grant and a National Science Foundation Postdoctoral Fellowship. Thanks to the organizers of the Algorithmic Number Theory program at MSRI, the Arizona Winter School in Tucson, and the Dwork Trimester in Padua for their hospitality, and to Laurent Berger, Pierre Colmez, Johan de Jong and the referee for helpful suggestions.

2 A few rings

In this chapter, we set some notations and conventions, and define some of the basic rings used in the local study of p-adic differential equations. We also review the basic properties of rings in which every finitely generated ideal is principal (Bézout rings), and introduce σ-modules and (σ, ∇)-modules.

2.1 Notations and conventions

Recall that for every field K of characteristic $p > 0$, there exists a complete discrete valuation ring with fraction field of characteristic 0, maximal ideal generated by p, and residue field isomorphic to K, and that this ring is unique up to noncanonical isomorphism. Such a ring is called a *Cohen ring* for K; see [Bo] for the basic properties of such rings. If K is perfect, the Cohen ring is unique up to canonical isomorphism, and coincides with the ring $W(K)$ of Witt vectors over K. (Note in passing: for K perfect, we use brackets to denote Teichmüller lifts into $W(K)$.)

Let k be a field of characteristic $p > 0$, and $C(k)$ a Cohen ring for k. Let $\mathcal{O}$ be a finite totally ramified extension of $C(k)$, let π be a uniformizer of $\mathcal{O}$, and fix once and for all a ring endomorphism σ_0 on $\mathcal{O}$ lifting the absolute Frobenius $x \mapsto x^p$ on k. Let $q = p^f$ be a power of p and put $\sigma = \sigma_0^f$. (In principle, one could dispense with σ_0 and simply take σ to be any ring endomorphism lifting the q-power Frobenius. We will eschew this additional level of generality so that we can invoke the results of [T1].) Let v_p denote the valuation on $\mathcal{O}[\frac{1}{p}]$ normalized so that $v_p(p) = 1$, and let $|\cdot|$ denote the norm on $\mathcal{O}[\frac{1}{p}]$ given by $|x| = p^{-v_p(x)}$.

Let $\mathcal{O}_0$ denote the fixed ring of $\mathcal{O}$ under σ. If k is algebraically closed, then the equation $u^\sigma = (\pi^\sigma/\pi)u$ in u has a nonzero solution modulo π, and so by a variant of Hensel's lemma (see Proposition 3.17) has a nonzero solution in $\mathcal{O}$. Then (π/u) is a uniformizer of $\mathcal{O}$ contained in $\mathcal{O}_0$, and hence $\mathcal{O}_0$ has the same value group as $\mathcal{O}$. That being the case, we can and will take $\pi \in \mathcal{O}_0$ in case k is algebraically closed.

We wish to alert the reader to several notational conventions in force throughout the paper. The first of these is "exponent consolidation". The expression $(x^{-1})^\sigma$, for x a ring element or matrix and σ a ring endomorphism, will often be abbreviated $x^{-\sigma}$. This is not to be confused with $x^{\sigma^{-1}}$; the former is the image under σ of the multiplicative inverse of x, the latter is the preimage of x under σ (if it exists). Similarly, if A is a matrix, then A^T will denote the transpose of A, and the expression $(A^{-1})^T$ will be abbreviated A^{-T}.

We will use the summation notation $\sum_{i=m}^n f(i)$ in some cases where $m > n$, in which case we mean 0 for $n = m - 1$ and $-\sum_{i=n+1}^{m-1} f(i)$ otherwise. The point of this convention is that $\sum_{i=m}^n f(i) = f(n) + \sum_{i=m}^{n-1} f(i)$ for all $n \in \mathbb{Z}$.

We will perform a number of calculations involving matrices; these will always be $n \times n$ matrices unless otherwise specified. Also, I will denote the $n \times n$ identity matrix over any ring.

2.2 Valued fields

Let $k((t))$ denote the field of Laurent series over k. Define a *valued field* to be an algebraic extension K of $k((t))$ for which there exist subextensions $k((t)) \subseteq L \subseteq M \subseteq N \subseteq K$ such that:

(a) $L = k^{1/p^m}((t))$ for some $m \in \{0, 1, \dots, \infty\}$;

(b) $M = k_M((t))$ for some separable algebraic extension $k_M/k^{1/p^m}$;

(c) $N = M^{1/p^n}$ for some $n \in \{0, 1, \ldots, \infty\}$;

(d) K is a separable totally ramified algebraic extension of N.

(Here F^{1/p^∞} means the perfection of the field F, and K/N totally ramified means that K and N have the same value group.) Note that n is uniquely determined by K: it is the largest integer n such that t has a p^n-th root in K. If $n < \infty$ (e.g., if $K/k((t))$ is finite), then L, M, N are also determined by K: k_M^{1/p^n} must be the integral closure of k in K, which determines k_M, and k^{1/p^m} must be the maximal purely inseparable subextension of k_M/k.

The following proposition shows that the definition of a valued field is only restrictive if k is imperfect. It also guides the construction of the rings Γ^K in Section 3.1.

Proposition 2.1. *If k is perfect, then any algebraic extension $K/k((t))$ is a valued field.*

Proof. Normalize the valuation v on $k((t))$ so that $v(t) = 1$. Let k_M be the integral closure of k in K, and define $L = k((t))$ and $M = k_M((t))$. Then (a) holds for $m = 0$ and (b) holds because k is perfect.

Let n be the largest nonnegative integer such that t has a p^n-th root in K, or ∞ if there is no largest integer. Put

$$N = \bigcup_{i=0}^{\infty} \left(K \cap M^{1/p^i} \right).$$

Since $t^{1/p^i} \in K$ for all $i \leq n$ and k_M is perfect, we have $M^{1/p^n} \subseteq N$. On the other hand, suppose $x^{1/p^i} \in (K \cap M^{1/p^i}) \setminus (K \cap M^{1/p^{i-1}})$, that is, $x \in M$ has a p^i-th root in K but has no p-th root in M. Then $v(x)$ is relatively prime to p, so we can find integers a and b such that $y = x^a/t^{bp^i} \in M$ has a p^i-th root in K and $v(y) = 1$. We can write every element of M uniquely as a power series in y, so every element of M has a p^i-th root in K. In particular, t has a p^i-th root in K, so $i \leq n$. We conclude that $N = M^{1/p^n}$, verifying (c).

If $y \in K^p \cap N$, then $y = z^p$ for some $z \in K$ and $y^{p^i} \in M$ for some i. Then $z^{p^{i+1}} \in M$, so $z \in N$. We thus have $K^p \cap N = N^p$, so K/N is separable. To verify that K/N is totally ramified, let K_0 be any finite subextension of $K/k((t))$ and let U be the maximal unramified subextension of $K_0/(K_0 \cap N)$. We now recall two basic facts from [Se] about finite extensions of fields complete with respect to discrete valuations:

1. K_0/U is totally ramified, because $K_0/(K_0 \cap N)$ and its residue field extension are both separable.

2. There is a unique unramified extension of $K_0 \cap N$ yielding any specified separable residue field extension.

Since $K_0 \cap N$ is a power series field, we can make unramified extensions of $K_0 \cap N$ with any specified residue field extension by extending the constant field $K_0 \cap k_M$. By the second assertion above, $U/(K_0 \cap N)$ must then be a constant field extension. However, k_M is integrally closed in K, so $U = K_0 \cap N$ and $K_0/(K_0 \cap N)$ is totally ramified by the first assertion above. Since K is the union of its finite subextensions over $k((t))$, we conclude that K/N is totally ramified, verifying (d). $\square$

The proposition fails for k imperfect, as there are separable extensions of $k((t))$ with inseparable residue field extensions. For example, if c has no p-th root in k, then $K = k((t))[x]/(x^p - x - ct^{-p})$ is separable over $k((t))$ but induces an inseparable residue field extension. Thus K cannot be a valued field, as valued fields contain their residue field extensions. What is true for any k is that valued fields finite and normal over $k((t))$ form a cofinal subset of the finite extensions of $k((t))$.

We denote the perfect and algebraic closures of $k((t))$ by $k((t))^{\mathrm{perf}}$ and $k((t))^{\mathrm{alg}}$; these are both valued fields. We denote the separable closure of $k((t))$ by $k((t))^{\mathrm{sep}}$; this is a valued field only if k is perfect, as we saw above.

We say a valued field K is *nearly separable* if it is a separable extension of $k^{1/p^i}((t))$ for some integer i. (That is, any inseparability is concentrated in the constant field.) This definition allows to approximate certain separability assertions for k perfect in the case of general k, where some separable extensions of K fail to be valued fields. For example,

$$k^{1/p}((t))[x]/(x^p - x - ct^{-p}) = k^{1/p}((t))[x]/((x - c^{1/p}t^{-1})^p - (x - c^{1/p}t^{-1}) - c^{1/p}t^{-1})$$

is a nearly separable valued field. In general, given any separable extension of $k((t))$, taking its compositum with $k^{1/p^i}((t))$ for sufficiently large i yields a a nearly separable valued field.

2.3 The "classical" case $K = k((t))$

The definitions and results of Chapter 3 generalize previously known definitions and results in the key case $K = k((t))$. We treat this case first, both to allow readers familiar with the prior constructions to get used to the notations of this paper, and to provide a base on which to build additional rings in Chapter 3.

For $K = k((t))$, let Γ^K be the ring of bidirectional power series $\sum_{i\in\mathbb{Z}} x_i u^i$, with $x_i \in \mathcal{O}$, such that $|x_i| \to 0$ as $i \to -\infty$. Note that Γ^K is a discrete valuation ring complete under the p-adic topology, whose residue field is isomorphic to K via the map $\sum x_i u^i \mapsto \sum \overline{x_i} t^i$ (using the bar to denote reduction modulo π). In particular, if $\pi = p$, then Γ^K is a Cohen ring for K.

For n in the value group of $\mathcal{O}$, we define the "naïve partial valuations"

$$v_n^{\mathrm{naive}}\left(\sum x_i u^i\right) = \min_{v_p(x_i)\leq n}\{i\},$$

taking the maximum to be $+\infty$ if no such i exist. These partial valuations obey some basic rules:

$$\begin{aligned} v_n(x+y) &\geq \min\{v_n(x), v_n(y)\} \\ v_n(xy) &\geq \min_m\{v_m(x) + v_{n-m}(y)\} \end{aligned}$$

In both cases, equality always holds if the minimum is achieved exactly once.

Define the *levelwise topology* on Γ^K by declaring the collection of sets

$$\{x \in \Gamma^K : v_n^{\mathrm{naive}}(x) > c\},$$

for each $c \in \mathbb{Q}$ and each n in the value group of $\mathcal{O}$, to be a neighborhood basis of 0. The levelwise topology is finer than the π-adic topology, and the Laurent polynomial ring $\mathcal{O}[u, u^{-1}]$ is dense in Γ^K under the levelwise topology; thus any levelwise continuous endomorphism of Γ^K is determined by the image of t.

The ring Γ^K_{con} is isomorphic to the subring of Γ^K consisting of those series $\sum_{i \in \mathbb{Z}} x_i u^i$ satisfying the more stringent convergence condition

$$\liminf_{i \to -\infty} \frac{v_p(x_i)}{-i} > 0.$$

It is also a discrete valuation ring with residue field K, but is not π-adically complete.

Using the naïve partial valuations, we can define actual valuations on certain subrings of Γ^K_{con}. Let $\Gamma^K_{r,\mathrm{naive}}$ be the set of $x = \sum x_i u^i$ in Γ^K_{con} such that $\lim_{n \to \infty} r v_n^{\mathrm{naive}}(x) + n = \infty$; the union of the subrings over all r is precisely Γ^K_{con}. (Warning: the rings $\Gamma^K_{r,\mathrm{naive}}$ for individual r are not discrete valuation rings, even though their union is.) On this subring, we have the function

$$w_r^{\mathrm{naive}}(x) = \min_n \{ r v_n^{\mathrm{naive}}(x) + n \} = \min_i \{ ri + v_p(x_i) \}$$

which can be seen to be a nonarchimedean valuation as follows. It is clear that $w_r^{\mathrm{naive}}(x + y) \geq \min\{w_r^{\mathrm{naive}}(x), w_r^{\mathrm{naive}}(y)\}$ from the inequality $v_n(x+y) \geq \min\{v_n(x), v_n(y)\}$. As for multiplication, it is equally clear that $w_r^{\mathrm{naive}}(xy) \geq w_r^{\mathrm{naive}}(x) + w_r^{\mathrm{naive}}(y)$; the subtle part is showing equality. Choose m and n minimal so that $w_r^{\mathrm{naive}}(x) = r v_m^{\mathrm{naive}}(x) + m$ and $w_r^{\mathrm{naive}}(y) = r v_n^{\mathrm{naive}}(y) + n$; then

$$r v_{m+n}^{\mathrm{naive}}(xy) + m + n \geq \min_i \{ r v_i^{\mathrm{naive}}(x) + i + r v_{m+n-i}^{\mathrm{naive}}(y) + m + n - i \}.$$

The minimum occurs only once, for $i = m$, so equality holds, yielding $w_r^{\mathrm{naive}}(xy) = w_r^{\mathrm{naive}}(x) + w_r^{\mathrm{naive}}(y)$.

Since w_r^{naive} is a valuation, we have a corresponding norm $|\cdot|_r^{\mathrm{naive}}$ given by $|x|_r^{\mathrm{naive}} = p^{-w_r^{\mathrm{naive}}(x)}$. This norm admits a geometric interpretation: the ring $\Gamma^K_{r,\mathrm{naive}}[\frac{1}{p}]$ consists of power series which converge and are bounded for $p^{-r} \leq |u| < 1$, where u runs over all finite extensions of $\mathcal{O}[\frac{1}{p}]$. Then $|\cdot|_r^{\mathrm{naive}}$ coincides with the supremum norm on the circle $|u| = p^{-r}$.

Recall that $\sigma_0 : \mathcal{O} \to \mathcal{O}$ is a lift of the p-power Frobenius on k. We choose an extension of σ_0 to a levelwise continuous endomorphism of Γ^K that maps Γ^K_{con} into itself, and which lifts the p-power Frobenius on $k((t))$. In other words, choose $y \in \Gamma^K_{\mathrm{con}}$ congruent to u^p modulo π, and define σ_0 by

$$\sum_i a_i u^i \mapsto \sum_i a_i^{\sigma_0} y^i.$$

Define $\sigma = \sigma_0^f$, where f is again given by $q = p^f$.

Let $\Gamma^K_{\mathrm{an,con}}$ be the ring of bidirectional power series $\sum_i x_i u^i$, now with $x_i \in \mathcal{O}[\frac{1}{p}]$, satisfying

$$\liminf_{i \to -\infty} \frac{v_p(x_i)}{-i} > 0, \qquad \liminf_{i \to +\infty} \frac{v_p(x_i)}{i} \geq 0.$$

In other words, for any series $\sum_i x_i u^i$ in $\Gamma^K_{\mathrm{an,con}}$, there exists $\eta > 0$ such that the series converges for $\eta \leq |u| < 1$. This ring is commonly known as the *Robba ring*. It contains $\Gamma^K_{\mathrm{con}}[\frac{1}{p}]$, as the subring of functions which are analytic and bounded on some annulus $\eta \leq |u| < 1$, but neither contains nor is contained in Γ^K.

We can view Γ^K as the π-adic completion of Γ^K_{con}; our next goal is to identify $\Gamma^K_{\mathrm{an,con}}$ with a certain completion of $\Gamma^K_{\mathrm{con}}[\frac{1}{p}]$. Let $\Gamma^K_{\mathrm{an},r,\mathrm{naive}}$ be the ring of series $x \in \Gamma^K_{\mathrm{an,con}}$ such that $rv_n^{\mathrm{naive}}(x) + n \to \infty$ as $n \to \pm\infty$. Then $\Gamma^K_{\mathrm{an,con}}$ is visibly the union of the rings $\Gamma^K_{\mathrm{an},r,\mathrm{naive}}$ over all $r > 0$. We equip $\Gamma^K_{\mathrm{an},r,\mathrm{naive}}$ with the Fréchet topology for the norms $|\cdot|_s^{\mathrm{naive}}$ for $0 < s \leq r$. These topologies are compatible with the embeddings $\Gamma^K_{\mathrm{an},r,\mathrm{naive}} \hookrightarrow \Gamma^K_{\mathrm{an},s,\mathrm{naive}}$ for $0 < s < r$ (that is, the topology on $\Gamma^K_{\mathrm{an},r,\mathrm{naive}}$ coincides with the subspace topology for the embedding), so by taking the direct limit we obtain a topology on $\Gamma^K_{\mathrm{an,con}}$, which by abuse of language we will also call the Fréchet topology. (A better name might be the "limit-of-Fréchet topology".) Note that $\Gamma^K_{r,\mathrm{naive}}[\frac{1}{p}]$ is dense in $\Gamma^K_{\mathrm{an},r,\mathrm{naive}}$ for each r, so $\Gamma^K_{\mathrm{con}}[\frac{1}{p}]$ is dense in $\Gamma^K_{\mathrm{an,con}}$.

Proposition 2.2. *The ring* $\Gamma^K_{\mathrm{an},r,\mathrm{naive}}$ *is complete (for the Fréchet topology).*

Proof. Let $\{x_i\}$ be a Cauchy sequence for the Fréchet topology, consisting of elements of $\Gamma^K_{r,\mathrm{naive}}[\frac{1}{p}]$. That means that for $0 < s \leq r$ and any $c > 0$, there exists N such that $w_s^{\mathrm{naive}}(x_i - x_j) > c$ for $i, j \geq N$. Write $x_i = \sum_l x_{i,l} u^l$; then for each l, $\{x_{i,l}\}$ forms a Cauchy sequence. More precisely, for $i, j \geq N$, we have

$$sl + v_p(x_{i,l} - x_{j,l}) > c.$$

Since $\mathcal{O}$ is complete, we can take the limit y_l of $\{x_{i,l}\}$, and it will satisfy $sl + v_p(x_{i,l} - y_l) > c$ for $i \geq N$. Thus if we can show $y = \sum_l y_l u^l \in \Gamma^K_{\mathrm{an},r,\mathrm{naive}}$, then $\{x_i\}$ will converge to y under $|\cdot|_s^{\mathrm{naive}}$ for each s.

Choose $s < r$ and $c > 0$; we must show that $sl + v_p(y_l) \geq c$ for all but finitely many l. There exists N such that $sl + v_p(x_{i,l} - y_l) \geq c$ for $i \geq N$. Choose a single such i; then

$$\begin{aligned} sl + v_p(y_l) &\geq \min\{sl + v_p(x_{i,l} - y_l), sl + v_p(x_{i,l})\} \\ &\geq \min\{c, sl + v_p(x_{i,l})\}. \end{aligned}$$

Since $x_i \in \Gamma^K_{r,\mathrm{naive}}[\frac{1}{p}]$, $sl + v_p(x_{i,l}) \geq c$ for all but finitely many l. For such l, we have $sl + v_p(y_l) \geq c$, as desired. Thus $y \in \Gamma^K_{\mathrm{an},r,\mathrm{naive}}$; as noted earlier, y is the limit of $\{x_i\}$ under each $|\cdot|_s^{\mathrm{naive}}$, and so is the Fréchet limit.

We conclude that each Cauchy sequence with elements in $\Gamma^K_{r,\mathrm{naive}}[\frac{1}{p}]$ has a limit in $\Gamma^K_{\mathrm{an},r,\mathrm{naive}}$. Since $\Gamma^K_{r,\mathrm{naive}}[\frac{1}{p}]$ is dense in $\Gamma^K_{\mathrm{an},r,\mathrm{naive}}$ (one sequence converging to $\sum_i x_i u^i$ is simply $\{\sum_{i \leq j} x_i u^i\}_{j=0}^{\infty}$), $\Gamma^K_{\mathrm{an},r,\mathrm{naive}}$ is complete for the Fréchet topology, as desired. □

Unlike Γ^K and Γ^K_{con}, $\Gamma^K_{\mathrm{an,con}}$ is not a discrete valuation ring. For one thing, π is invertible in $\Gamma^K_{\mathrm{an,con}}$. For another, there are plenty of noninvertible elements of $\Gamma^K_{\mathrm{an,con}}$, such as

$$\prod_{i=1}^{\infty} \left(1 - \frac{u^{p^i}}{p^i}\right).$$

For a third, $\Gamma^K_{\mathrm{an,con}}$ is not Noetherian; the ideal $(x_1, x_2, \dots)$, where

$$x_j = \prod_{i=j}^{\infty} \left(1 - \frac{u^{p^i}}{p^i}\right),$$

is not finitely generated. However, as long as we restrict to "finite" objects, $\Gamma^K_{\mathrm{an,con}}$ behaves well: a theorem of Lazard [L] (see also [Cr2, Proposition 4.6] and our own Section 3.6) states that $\Gamma^K_{\mathrm{an,con}}$ is a *Bézout ring*, which is to say every finitely generated ideal is principal.

For L a finite extension of $k((t))$, we have $L \cong k'((t'))$ for some finite extension k' of k and some uniformizer t', so one could define Γ^L, Γ^L_{con}, $\Gamma^L_{\mathrm{an,con}}$ abstractly as above. However, a better strategy will be to construct these in a "relative" fashion; the results will be the same as the abstract rings, but the relative construction will give us more functoriality, and will allow us to define $\Gamma^L, \Gamma^L_{\mathrm{con}}, \Gamma^L_{\mathrm{an,con}}$ even when L is an infinite algebraic extension of $k((t))$. We return to this approach in Chapter 3.

The rings defined above occur in numerous other contexts, so it is perhaps not surprising that there are several sets of notation for them in the literature. One common set is

$$\mathcal{E} = \Gamma^{k((t))}[\tfrac{1}{p}], \qquad \mathcal{E}^\dagger = \Gamma^{k((t))}_{\mathrm{con}}[\tfrac{1}{p}], \qquad \mathcal{R} = \Gamma^{k((t))}_{\mathrm{an,con}}.$$

The peculiar-looking notation we have set up will make it easier to deal systematically with a number of additional rings we will be defining in Chapter 3.

2.4 More on Bézout rings

Since $\Gamma^K_{\mathrm{an,con}}$ is a Bézout ring, as are trivially all discrete valuation rings, it will be useful to record some consequences of the Bézout property.

Lemma 2.3. *Let R be a Bézout ring. If $x_1, \dots, x_n \in R$ generate the unit ideal, then there exists a matrix A over R with determinant 1 such that $A_{1i} = x_i$ for $i = 1, \dots, n$.*

Proof. We prove this by induction on n, the case $n = 1$ being evident. Let d be a generator of $(x_1, \dots, x_{n-1})$. By the induction hypothesis, we can find an $(n-1) \times (n-1)$ matrix B of determinant 1 such that $B_{1i} = x_i/d$ for $i = 1, \dots, n-1$; extend B to an $n \times n$ matrix by setting $B_{nn} = 1$ and $B_{in} = B_{ni} = 0$ for $i = 1, \dots, n-1$. Since $(d, x_n) = (x_1, \dots, x_n)$ is the unit ideal, we can find $e, f \in R$ such that $de - fx_n = 1$. Define the matrix

$$C = \begin{pmatrix} d & 0 & \cdots & 0 & x_n \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \cdots & 1 & 0 \\ f & 0 & \cdots & 0 & e \end{pmatrix}; \qquad \text{that is,} \qquad C_{ij} = \begin{cases} d & i = j = 1 \\ 1 & 2 \leq i = j \leq n-1 \\ e & i = j = n \\ x_n & i = 1, j = n \\ f & i = n, j = 1 \\ 0 & \text{otherwise.} \end{cases}$$

Then we may take $A = CB$. ☐

Given a finite free module M over a domain R, we may regard M as a subset of $M \otimes_R \operatorname{Frac}(R)$; given a subset S of M, we define the *saturated span* $\operatorname{SatSpan}(S)$ of S as the intersection of M with the $\operatorname{Frac}(R)$-span of S within $M \otimes_R \operatorname{Frac}(R)$. Note that the following lemma does not require any finiteness condition on S.

Lemma 2.4. *Let M be a finite free module over a Bézout domain R. Then for any subset S of M, $\operatorname{SatSpan}(S)$ is free and admits a basis that extends to a basis of M; in particular, $\operatorname{SatSpan}(S)$ is a direct summand of M.*

Proof. We proceed by induction on the rank of M, the case of rank 0 being trivial. Choose a basis $\mathbf{e}_1, \dots, \mathbf{e}_n$ of M. If S is empty, there is nothing to prove; otherwise, choose $\mathbf{v} \in S$ and write $\mathbf{v} = \sum_i c_i \mathbf{e}_i$. Since R is a Bézout ring, we can choose a generator r of the ideal $(c_1, \dots, c_n)$. Put $\mathbf{w} = \sum_i (c_i/r)\mathbf{e}_i$; then $\mathbf{w} \in \operatorname{SatSpan}(S)$ since $r\mathbf{w} = \mathbf{v}$. By Lemma 2.3, there exists an invertible matrix A over R with $A_{1i} = c_i/r$. Put $\mathbf{y}_j = \sum_i A_{ji}\mathbf{e}_i$ for $j = 2, \dots, n$; then $\mathbf{w}$ and the $\mathbf{y}_j$ form a basis of M (because A is invertible), so $M/\operatorname{SatSpan}(\mathbf{w})$ is free. Thus the induction hypothesis applies to $M/\operatorname{SatSpan}(\mathbf{w})$, where the saturated span of the image of S admits a basis $\mathbf{x}_1, \dots, \mathbf{x}_r$. Together with $\mathbf{w}$, any lifts of $\mathbf{x}_1, \dots, \mathbf{x}_r$ to M form a basis of $\operatorname{SatSpan}(S)$ that extends to a basis of M, as desired. $\square$

Note that the previous lemma immediately implies that every finite torsion-free module over R is free. (If M is torsion-free and $\phi : F \to M$ is a surjection from a free module F, then $\ker(\phi)$ is saturated, so $M \cong F/\ker(\phi)$ is free.) A similar argument yields the following vitally important fact.

Proposition 2.5. *Let R be a Bézout domain. Then every finite locally free module over R is free.*

Proof. Let M be a finite locally free module over R. Since $\operatorname{Spec} R$ is connected, the localizations of M all have the same rank r. Choose a surjection $\phi : F \to M$, where F is a finite free R-module, and let $N = \operatorname{SatSpan}(\ker(\phi))$. Then we have a surjection $M \cong F/\ker(\phi) \to F/N$, and F/N is free. Tensoring ϕ with $\operatorname{Frac}(R)$, we obtain a surjection $F \otimes_R \operatorname{Frac}(R) \to M \otimes_R \operatorname{Frac}(R)$ of vector spaces of dimensions n and r. Thus the kernel of this map has rank $n - r$, which implies that N has rank $n - r$ and F/N is free of rank r.

Now localizing at each prime $\mathfrak{p}$ of R, we obtain a surjection $M_{\mathfrak{p}} \to (F/N)_{\mathfrak{p}}$ of free modules of the same rank. By a standard result, this map is in fact a bijection. Thus $M \to F/N$ is locally bijective, hence is bijective, and M is free as desired. $\square$

The following lemma is a weak form of Galois descent for Bézout rings; its key value is that it does not require that the ring extension be finite.

Lemma 2.6. *Let R_1/R_2 be an extension of Bézout domains and G a group of automorphisms of R_1 over R_2, with fixed ring R_2. Assume that every G-stable finitely generated ideal of R_1 contains a nonzero element of R_2. Let M_2 be a finite free module over R_2 and N_1 a saturated G-stable submodule of $M_1 = M_2 \otimes_{R_2} R_1$ stable under G. Then N_1 is equal to $N_2 \otimes_{R_2} R_1$ for a saturated submodule (necessarily unique) N_2 of M_2.*

Proof. We induct on $n = \operatorname{rank} M_2$, the case $n = 0$ being trivial. Let $\mathbf{e}_1, \dots, \mathbf{e}_n$ be a basis of M_2, and let P_1 be the intersection of N_1 with the span of $\mathbf{e}_2, \dots, \mathbf{e}_n$; since N_1 is saturated, P_1 is a direct summand of $\operatorname{SatSpan}(\mathbf{e}_2, \dots, \mathbf{e}_n)$ by Lemma 2.4 and hence also of M_1. By the induction hypothesis, $P_1 = P_2 \otimes_{R_2} R_1$ for a saturated submodule P_2 of M_2 (necessarily a direct summand by Lemma 2.4). If $N_1 = P_1$, we are done. Otherwise, N_1/P_1 is a G-stable, finitely generated ideal of R_1 (since N_1 is finitely generated by Lemma 2.4), so contains a nonzero element c of R_2. Pick $\mathbf{v} \in N_1$ reducing to c; that is, $\mathbf{v} - c\mathbf{e}_1 \in \operatorname{SatSpan}(\mathbf{e}_2, \dots, \mathbf{e}_n)$.

Pick generators $\mathbf{w}_1, \dots, \mathbf{w}_m$ of P_2; since P_2 is a direct summand of $\operatorname{SatSpan}(\mathbf{e}_2, \dots, \mathbf{e}_n)$, we can choose $\mathbf{x}_1, \dots, \mathbf{x}_{n-m-1}$ in M_2 so that $\mathbf{e}_1, \mathbf{w}_1, \dots, \mathbf{w}_m, \mathbf{x}_1, \dots, \mathbf{x}_{n-m-1}$ is a basis of M. In this basis, we may write $\mathbf{v} = c\mathbf{e}_1 + \sum_i d_i \mathbf{w}_i + \sum_i f_i \mathbf{x}_i$, where c is the element of R_2 chosen above. Put $\mathbf{y} = \mathbf{v} - \sum_i d_i \mathbf{w}_i$. For any $\tau \in G$, we have $\mathbf{y}^\tau = c\mathbf{e}_1 + \sum_i f_i^\tau \mathbf{x}_i$, so on one hand, $\mathbf{y}^\tau - \mathbf{y}$ is a linear combination of $\mathbf{x}_1, \dots, \mathbf{x}_{n-m-1}$. On the other hand, $\mathbf{y}^\tau - \mathbf{y}$ belongs to N_1 and so is a linear combination of $\mathbf{w}_1, \dots, \mathbf{w}_m$. This forces $\mathbf{y}^\tau - \mathbf{y} = 0$ for all $\tau \in G$; since G has fixed ring R_2, we conclude $\mathbf{y}$ is defined over R_2. Thus we may take $N_2 = \operatorname{SatSpan}(\mathbf{y}, \mathbf{w}_1, \dots, \mathbf{w}_m)$. ☐

Note that the hypothesis that every G-stable finitely generated ideal of R_1 contains a nonzero element of R_2 is always satisfied if G is finite: for any nonzero r in the ideal, $\prod_{\tau \in G} r^\tau$ is nonzero and G-stable, so belongs to R_2.

2.5 σ-modules and (σ, ∇)-modules

The basic object in the local study of p-adic differential equations is a module with connection and Frobenius structure. In our approach, we separate these two structures and study the Frobenius structure closely before linking it with the connection. To this end, in this section we introduce σ-modules and (σ, ∇)-modules, and outline some basic facts of what might be dubbed "semilinear algebra". These foundations, in part, date back to Katz [Ka] and were expanded by de Jong [dJ].

For R an integral domain in which $p \neq 0$ and σ a ring endomorphism of R, we define a *σ-module* over R to be a finite locally free R-module M equipped with an R-linear map $F : M \otimes_{R,\sigma} R \to M$ that becomes an isomorphism over $R[\frac{1}{p}]$; the tensor product notation indicates that R is viewed as an R-module via σ. Note that we will only use this definition when R is a Bézout ring, in which case every finite locally free R-module is actually free by Proposition 2.5. Then to specify F, it is equivalent to specify an additive, σ-linear map from M to M that acts on any basis of M by a matrix invertible over $R[\frac{1}{p}]$. We abuse notation and refer to this map as F as well; since we will only use the σ-linear map in what follows (with one exception: in proving Proposition 6.11), there should not be any confusion induced by this.

Now suppose R is one of $\Gamma^K, \Gamma^K[\frac{1}{p}], \Gamma^K_{\mathrm{con}}, \Gamma^K_{\mathrm{con}}[\frac{1}{p}]$ or $\Gamma^K_{\mathrm{an,con}}$ for $K = k((t))$. Let Ω^1_R be the free module over R generated by a single symbol du, and let $d : R \to \Omega^1_R$ be the $\mathcal{O}$-linear derivation given by the formula

$$d\left(\sum_i x_i u^i\right) = \sum_i i x_i u^{i-1}\,du.$$

We define a (σ, ∇)-module over R to be a σ-module M plus a connection $\nabla : M \to M \otimes_R \Omega^1_R$ (i.e., an additive map satisfying the Leibniz rule $\nabla(c\mathbf{v}) = c\nabla(\mathbf{v}) + \mathbf{v} \otimes dc$ for $c \in R$ and $\mathbf{v} \in M$) that makes the following diagram commute:

$$\begin{tikzcd} M \arrow[r, "\nabla"] \arrow[d, "F"] & M \otimes \Omega^1_R \arrow[d, "F \otimes d\sigma"] \\ M \arrow[r, "\nabla"] & M \otimes \Omega^1_R \end{tikzcd}$$

Warning: this definition is not the correct one in general. For larger rings R, one must include the condition that ∇ is integrable. That is, writing ∇_1 for the induced map $M \otimes_R \Omega^1_R \to M \otimes_R \wedge^2 \Omega^1_R$, we have $\nabla_1 \circ \nabla = 0$. This condition is superfluous in our context because Ω^1_R has rank one, so ∇_1 is automatically zero.

A *morphism* of σ-modules or (σ, ∇)-modules is a homomorphism of the underlying R-modules compatible with the additional structure in the obvious fashion. An *isomorphism* of σ-modules or (σ, ∇)-modules is a morphism admitting an inverse; an *isogeny* is a morphism that becomes an isomorphism over $R[\frac{1}{p}]$.

Direct sums, tensor products, exterior powers, and subobjects are defined in the obvious fashion, as are duals if $p^{-1} \in R$; quotients also make sense provided that the quotient R-module is locally free. In particular, if $M_1 \subseteq M_2$ is an inclusion of σ-modules, the saturation of M_1 in M_2 is also a σ-submodule of M_1; if M_1 itself is saturated, the quotient M_2/M_1 is locally free and hence is a σ-module.

Given λ fixed by σ, we define the *twist* of a σ-module M by λ as the σ-module with the same underlying module but whose Frobenius has been multiplied by λ.

We say a σ-module M is *standard* if it is isogenous to a σ-module with a basis $\mathbf{v}_1, \dots, \mathbf{v}_n$ such that $F\mathbf{v}_i = \mathbf{v}_{i+1}$ for $i = 1, \dots, n-1$ and $F\mathbf{v}_n = \lambda\mathbf{v}_1$ for some $\lambda \in R$ fixed by σ. (The fact that λ is fixed by σ is inserted for convenience only.) If M is actually a (σ, ∇)-module, we say M is standard as a (σ, ∇)-module if the same condition holds with the additional restriction that $\nabla\mathbf{v}_i = 0$ for $i = 1, \dots, n$ (i.e., the $\mathbf{v}_i$ are "horizontal sections" for the connection). If $\mathbf{v}$ is a nonzero element of M such that $F\mathbf{v} = \lambda\mathbf{v}$ for some λ, we say $\mathbf{v}$ is an *eigenvector* of M of *eigenvalue* λ and *slope* $v_p(\lambda)$.

Warning: elsewhere in the literature, the slope may be normalized differently, namely as $v_p(\lambda)/v_p(q)$. (Recall that $q = p^f$.) Since we will hold q fixed, this normalization will not affect our results.

From Lemma 2.6, we have the following descent lemma for σ-modules. (The condition on G-stable ideals is satisfied because R_1/R_2 is an unramified extension of discrete valuation rings.)

Corollary 2.7. *Let R_1/R_2 be an unramified extension of discrete valuation rings unramified over $\mathcal{O}$, and let σ be a ring endomorphism of R_1 carrying R_2 into itself. Let* $\mathrm{Gal}^\sigma(R_1/R_2)$ *be the group of automorphisms of R_1 over R_2 commuting with σ; assume that this group has fixed ring R_2. Let M_2 be a σ-module over R_2 and N_1 a saturated σ-submodule of $M_1 = M_2 \otimes_{R_2} R_1$ stable under* $\mathrm{Gal}^\sigma(R_1/R_2)$. *Then $N_1 = N_2 \otimes_{R_2} R_1$ for some σ-submodule N_2 of M_2.*

3 A few more rings

In this chapter, we define a number of additional auxiliary rings used in our study of σ-modules. Again, we advise the reader to skim this chapter on first reading and return to it as needed.

3.1 Cohen rings

We proceed to generalizing the constructions of Section 2.3 to valued fields. This cannot be accomplished using Witt vectors because $k((t))$ and its finite extensions are not perfect. To get around this, we fix once and for all a levelwise continuous Frobenius lift σ_0 on $\Gamma^{k((t))}$ carrying $\Gamma^{k((t))}_{\mathrm{con}}$ into itself; all of our constructions will be made relative to the choice of σ_0.

Recall that a valued field K is defined to be an algebraic extension of $k((t))$ admitting subextensions $k((t)) \subseteq L \subseteq M \subseteq N \subseteq K$ such that:

(a) $L = k^{1/p^m}((t))$ for some $m \in \{0, 1, \dots, \infty\}$;

(b) $M = k_M((t))$ for some separable algebraic extension $k_M/k^{1/p^m}$;

(c) $N = M^{1/p^n}$ for some $n \in \{0, 1, \dots, \infty\}$;

(d) K is a separable totally ramified algebraic extension of N.

We will associate to each valued field K a complete discrete valuation ring Γ^K unramified over $\mathcal{O}$, equipped with a Frobenius lift σ_0 extending the definition of σ_0 on $\Gamma^{k((t))}$. This assignment will be functorial in K.

Let $\mathcal{C}$ be the category of complete discrete valuation rings unramified over $\mathcal{O}$, in which morphisms are unramified morphisms of rings (i.e., morphisms which induce isomorphisms of the value groups). If $R_0, R_1 \in \mathcal{C}$ have residue fields k_0, k_1 and a homomorphism $\phi : k_0 \to k_1$ is given, we say the morphism $f : R_0 \to R_1$ is *compatible* (with ϕ) if the diagram

$$\begin{tikzcd}
R_0 \arrow[r, "f"] \arrow[d] & R_1 \arrow[d] \\
k_0 \arrow[r, "\phi"] & k_1
\end{tikzcd}$$

commutes.

Lemma 3.1. *Let k_1/k_0 be a finite separable extension of fields, and take $R_0 \in \mathcal{C}$ with residue field k_0. Then there exists $R_1 \in \mathcal{C}$ with residue field k_1 and a compatible morphism $R_0 \to R_1$.*

Proof. By the primitive element theorem, there exists a monic separable polynomial $\overline{P}(x)$ over k_0 and an isomorphism $k_1 \cong k_0[x]/(\overline{P}(x))$. Choose a monic polynomial $P(x)$ over R_0 lifting $\overline{P}(x)$ and set $R_1 = R_0[x]/(P(x))$. Then the inclusion $R_0 \to R_0[x]$ induces the desired morphism $R_0 \to R_1$. ☐

Lemma 3.2. *Let $k_0 \to k_1 \to k_2$ be homomorphisms of fields, with k_1/k_0 finite separable. For $i = 0, 1, 2$, take $R_i \in \mathcal{C}$ with residue field k_i. Let $f : R_0 \to R_1$ and $g : R_0 \to R_2$ be compatible morphisms. Then there exists a unique compatible morphism $h : R_1 \to R_2$ such that $g = h \circ f$.*

Proof. As in the previous proof, choose a monic separable polynomial $\overline{P}(x)$ over k_0 and an isomorphism $k_1 \cong k_0[x]/(\overline{P}(x))$. Let y be the image of $x + (\overline{P}(x))$ in k_1, and let z be the image of y in k_2.

Choose a monic polynomial $P(x)$ over R_0 lifting $\overline{P}(x)$, and view R_0 as a subring of R_1 and R_2 via f and g, respectively. By Hensel's lemma, there exist unique roots α and β of $P(x)$ in R_1 and R_2 reducing to y and z, respectively, so h must satisfy $h(\alpha) = \beta$ if it exists. Then $R_0[x]/(P(x)) \cong R_1$ by the map sending x to α and $R_0[x]/(P(x)) \hookrightarrow R_2$ by the map sending x to β, so there exists a unique $h : R_1 \to R_2$ such that $h(\alpha) = \beta$, and this gives the desired morphism. ☐

Corollary 3.3. *If k_1/k_0 is finite Galois, and $R_i \in \mathcal{C}$ has residue field k_i for $i = 0, 1$, then for any compatible morphism $f : R_0 \to R_1$, the group of f-equivariant automorphisms of R_1 is isomorphic to* $\mathrm{Gal}(k_1/k_0)$.

Proof. Apply Lemma 3.2 with $k_0 \to k_1$ the given embedding and $k_1 \to k_1$ an element of $\mathrm{Gal}(k_1/k_0)$; the resulting h is the corresponding automorphism. ☐

Corollary 3.4. *If k_1/k_0 is finite separable, $R_i \in \mathcal{C}$ has residue field k_i for $i = 0, 1$, and $f : R_0 \to R_1$ is a compatible morphism, then any compatible endomorphism of R_0 admits a unique f-equivariant extension to R_1.*

Proof. If $\phi : R_0 \to R_0$ is the given endomorphism, apply Lemma 3.2 with $g = f \circ \phi$. ☐

For m a nonnegative integer, let $\mathcal{O}_m$ be a copy of $\mathcal{O}$. Then the assignment $k^{1/p^m} \rightsquigarrow \mathcal{O}_m$ is functorial via the morphism σ_0^i compatible with $k^{1/p^m} \to k^{1/p^{m+i}}$; thus we can define $\mathcal{O}_\infty$ as the completed direct limit of the $\mathcal{O}_m$. For any finite separable extension k_M of k^{1/p^m}, choose $\mathcal{O}_M$ in $\mathcal{C}$ according to Lemma 3.1, to obtain a compatible morphism $\mathcal{O}_m \to \mathcal{O}_M$; note that $\mathcal{O}_M$ is unique up to *canonical* isomorphism by Lemma 3.2. Moreover, this assignment is functorial in k_M (again by Lemma 3.2), so again we may pass to infinite extensions by taking the completed direct limit.

Now suppose K is a valued field finite over $k((t))$ and L, M, k_M, N, n are as in the definition of valued fields; recall that since $n < \infty$, L, M, k_M, N are uniquely determined by

K. Define $\mathcal{O}_M$ associated to k_M as above, define Γ^M as the ring of power series $\sum_{i \in \mathbb{Z}} a_i u^i$, with $a_i \in \mathcal{O}_M$, such that $|a_i| \to 0$ as $i \to -\infty$, and identify $\Gamma^M / \pi \Gamma^M$ with $M = k_M((t))$ via the map $\sum_i a_i u^i \mapsto \sum_i \overline{a_i} t^i$. Define Γ^N as a copy of Γ^M, but with Γ^M embedded via σ_0^n (which makes sense since $n < \infty$), and identify the residue field of Γ^N with N compatibly. Define Γ^K as a copy of Γ^N with its residue field identified with K via some continuous k_M^{1/p^n}-algebra isomorphism $K \cong N$ (which exists because both fields are power series fields over k_M^{1/p^n} by the Cohen structure theorem). Once this choice is made, there is a unique (necessarily levelwise continuous) morphism $\Gamma^N \to \Gamma^K$ compatible with the embedding $N \hookrightarrow K$. The assignments of $\Gamma^M, \Gamma^N, \Gamma^K$ are functorial, again by Lemma 3.2, so again we may extend the definition to infinite K by completion.

Note that if $K/k((t))$ is finite, then Γ^K is equipped with a levelwise topology, and the embeddings provided by functoriality are levelwise continuous. Moreover, σ_0 extends uniquely to each Γ^K, and the functorial morphisms are σ_0-equivariant.

If k and K are perfect and $\mathcal{O} = C(k) = W(k)$, then Γ^K is canonically isomorphic to the Witt ring $W(K)$. Under that isomorphism, σ_0 corresponds to the Witt vector Frobenius, which sends each Teichmüller lift to its p-th power. For general $\mathcal{O}$, we have $\Gamma^K \cong W(K) \otimes_{W(k)} \mathcal{O}$.

We will often fix a field K (typically $k((t))$ itself) and write Γ instead of Γ^K. In this case, we will frequently refer to Γ^L for various canonical extensions L of K, such as the separable closure K^{sep}, the perfect closure K^{perf}, and the algebraic closure K^{alg}. In all of these cases, we will drop the K from the notation where it is understood, writing Γ^{perf} for $\Gamma^{K^{\mathrm{perf}}}$ and so forth.

3.2 Overconvergent rings

Let K be a valued field. Let v_K denote the valuation on K extending the valuation on $k((t))$, normalized so that $v_K(t) = 1$. Again, let $q = p^f$, and put $\sigma = \sigma_0^f$ on Γ^K. We define a subring Γ^K_{con} of Γ^K of "overconvergent" elements; the construction will not look quite like the construction of $\Gamma^{k((t))}_{\mathrm{con}}$ from Section 2.3, so we must check that the two are consistent.

First assume K is perfect. For $x \in \Gamma^K[\frac{1}{p}]$, write $x = \sum_{i=m}^{\infty} \pi^i [\overline{x_i}]$, where $m v_p(\pi) = v_p(x)$, each $\overline{x_i}$ belongs to K and the brackets denote Teichmüller lifts. For n in the value group of $\mathcal{O}$, we define the "partial valuations"

$$v_n(x) = \min_{v_p(\pi^i) \leq n} \{ v_K(\overline{x_i}) \}.$$

These partial valuations obey two rules analogous to those for their naïve counterparts, plus a third that has no analogue:

$$\begin{aligned} v_n(x+y) &\geq \min\{v_n(x), v_n(y)\} \\ v_n(xy) &\geq \min_m \{v_m(x) + v_{n-m}(y)\} \\ v_n(x^\sigma) &= q v_n(x). \end{aligned}$$

Again, equality holds in the first two lines if the minimum is achieved exactly once.

For each $r > 0$, let Γ_r^K denote the subring of $x \in \Gamma^K$ such that $\lim_{n\to\infty}(rv_n(x)+n) = \infty$. On $\Gamma_r^K[\frac{1}{p}] \setminus \{0\}$, we define the function

$$w_r(x) = \min_n \{rv_n(x) + n\};$$

then w_r is a nonarchimedean valuation by the same argument as for w_r^{naive} given in Section 2.3. Define $\Gamma_{\text{con}}^K = \cup_{r>0}\Gamma_r^K$.

The rings Γ_r^K will be quite useful, but one must handle them with some caution, for the following reasons:

(a) The map $\sigma : \Gamma^K \to \Gamma^K$ sends Γ_{con}^K into itself, but does not send Γ_r^K into itself; rather, it sends Γ_r^K into $\Gamma_{r/q}^K$.

(b) The ring Γ_{con}^K is a discrete valuation ring, but the rings Γ_r^K are not.

(c) The ring Γ_r^K is complete for w_r, but not for the p-adic valuation.

For K arbitrary, we want to define Γ_{con}^K as $\Gamma_{\text{con}}^{\text{alg}} \cap \Gamma^K$. This intersection is indeed a discrete valuation ring (so again its fraction field is obtained by adjoining $\frac{1}{p}$), but it is not clear that its residue field is all of K. Indeed, it is *a priori* possible that the intersection is no larger than $\mathcal{O}$ itself! In fact, this pathology does not occur, as we will see below.

To make that definition, we must also check that $\Gamma_{\text{con}}^{\text{alg}} \cap \Gamma^{k((t))}$ coincides with the ring $\Gamma_{\text{con}}^{k((t))}$ defined earlier. This is obvious in a special case: if $\sigma_0(u) = u^p$, then u is a Teichmüller lift in $\Gamma_{\text{con}}^{\text{alg}}$, and in this case one can check that the partial valuations and naïve partial valuations coincide. In general they do not coincide, but in a sense they are not too far apart. The relationship might be likened to that between the naïve and canonical heights on an abelian variety over a number field.

Put $z = u^\sigma/u^q - 1$. By the original definition of σ on $\Gamma^{k((t))}$, $v_p(z) > 0$ and $z \in \Gamma_{\text{con}}^{k((t))}$. That means we can find $r > 0$ such that $rv_n^{\text{naive}}(z) + n > 0$ for all n; for all $s \leq r$, we then have $w_s^{\text{naive}}(u^\sigma/u^q) = 0$.

Lemma 3.5. *Choose $r > 0$ such that $rv_n^{\text{naive}}(z)+n > 0$ for all n. For $x = \sum_i x_i u^i$ in $\Gamma_{r,\text{naive}}^{k((t))}$, if $0 < s \leq qr$ and $w_s^{\text{naive}}(x) \geq c$, then $w_{s/q}^{\text{naive}}(x^\sigma) \geq c$.*

Proof. We have

$$\begin{aligned} w_{s/q}^{\text{naive}}(x_i^\sigma (u^i)^\sigma) &= w_{s/q}^{\text{naive}}(x_i u^{qi} (u^\sigma/u^q)^i) \\ &= w_{s/q}^{\text{naive}}(x_i u^{qi}) + w_{s/q}^{\text{naive}}((u^\sigma/u^q)^i) \\ &= w_s^{\text{naive}}(x_i u^i) \end{aligned}$$

since $w_{s/q}^{\text{naive}}(u^\sigma/u^q) = 0$ whenever $s/q \leq r$.

Given that $w_s^{\text{naive}}(x) \geq c$, it follows that $w_{s/q}^{\text{naive}}(x_i u^i) \geq c$ for each i, and by the above argument, that $w_{s/q}^{\text{naive}}(x_i^\sigma (u^i)^\sigma) \geq c$. We conclude that $w_{s/q}^{\text{naive}}(x^\sigma) \geq c$, as desired. ☐

Lemma 3.6. *Choose $r > 0$ such that $rv_n^{\mathrm{naive}}(z) + n > 0$ for all n. For $x \in \Gamma_{r,\mathrm{naive}}^{k((t))}$, if $sv_j^{\mathrm{naive}}(x) + j \geq c$ for all $j \leq n$, then $sv_j(x) + j \geq c$ for all $j \leq n$.*

Proof. Note that $v_0 = v_0^{\mathrm{naive}}$, so the desired result holds for $n = 0$; we prove the general result by induction on n. Suppose, as the induction hypothesis, that if $sv_j^{\mathrm{naive}}(x) + j \geq c$ for all $j < n$, then $sv_j(x) + j \geq c$ for all $j < n$. Before deducing the desired result, we first study the special case $x = u$ in detail (but using the induction hypothesis in full generality).

Choose i large enough that

$$v_p([t] - (u^{\sigma^{-i}})^{q^i}) > n.$$

Then

$$\begin{aligned} v_n(u) &\geq \min\{v_n([t]), v_n(u - [t])\} \\ &= \min\{1, v_n(u - (u^{\sigma^{-i}})^{q^i})\}. \end{aligned}$$

Applying σ^i yields

$$q^i v_n(u) \geq \min\{q^i, v_n(u^{\sigma^i} - u^{q^i})\}.$$

Since $u \in \Gamma_{r,\mathrm{naive}}^{k((t))}$ and $w_r^{\mathrm{naive}}(u) = r$ trivially, we may apply Lemma 3.5 to $u, u^\sigma, \dots, u^{\sigma^{i-1}}$ in succession to obtain

$$w_{r/q^i}^{\mathrm{naive}}(u^{\sigma^i}) \geq r.$$

Since $w_{r/q^i}^{\mathrm{naive}}(u^{q^i}) = r$, we conclude that $w_{r/q^i}^{\mathrm{naive}}(u^{\sigma^i} - u^{q^i}) \geq r$.

Let $y = (u^{\sigma^i} - u^{q^i})/\pi$. Then for $j \leq n - v_p(\pi)$,

$$\begin{aligned} (r/q^i)v_j^{\mathrm{naive}}(y) + j &= (r/q^i)v_{j+v_p(\pi)}^{\mathrm{naive}}(y\pi) + j + v_p(\pi) - v_p(\pi) \\ &\geq w_{r/q^i}^{\mathrm{naive}}(y\pi) - v_p(\pi) \\ &\geq r - v_p(\pi). \end{aligned}$$

By the induction hypothesis, we conclude that $(r/q^i)v_{n-v_p(\pi)}(y) + n - v_p(\pi) \geq r - v_p(\pi)$, and so $(r/q^i)v_n(y\pi) + n \geq r$. From above, we have

$$\begin{aligned} q^i v_n(u) &\geq \min\{q^i, v_n(u^{\sigma^i} - u^{q^i})\} \\ &\geq \min\{q^i, q^i - q^i n/r\} \\ &= q^i - q^i n/r. \end{aligned}$$

Thus $rv_n(u) + n \geq r$. Since $v_n(u) \leq 1$, we also have $sv_n(u) + n \geq s$ for $s \leq r$; that is, the desired conclusion holds for the special case $x = u$. By the multiplication rule for partial valuations, we also have $sv_n(u^i) + n \geq si$ for all i.

With the case $x = u$ in hand, we now prove the desired conclusion for general x. We are given $sv_j^{\mathrm{naive}}(x) + j \geq c$ for $j \leq n$; by the induction hypothesis, all that we must prove is that $sv_n(x) + n \geq c$.

The assumption $sv_j^{\mathrm{naive}}(x) + j \geq c$ implies that $sv_j^{\mathrm{naive}}(x_i u^i) + j \geq c$ for all $j \leq n$, which is to say, if $v_p(x_i) \leq n$ then $si + v_p(x_i) \geq c$. For $j = v_p(x_i)$, we have

$$\begin{aligned} sv_n(x_i u^i) + n &= sv_{n-j}(u^i) + n - j + j \\ &\geq si + j \\ &\geq c. \end{aligned}$$

We conclude that $sv_n(x) + n \geq c$, completing the induction. $\square$

We next refine the previous result as follows.

Lemma 3.7. *Choose $r > 0$ such that $rv_n^{\mathrm{naive}}(z) + n > 0$ for all n. If $x \in \Gamma^{k((t))}_{r,\mathrm{naive}}$, then for any $s \leq r$, $\min_{j \leq n}\{sv_j^{\mathrm{naive}}(x) + j\} = \min_{j \leq n}\{sv_j(x) + j\}$ for all n. In particular, $w_s^{\mathrm{naive}}(x) = w_s(x)$.*

That is, the naïve valuations w_s^{naive} are not so simple-minded after all; as long as s is not too large, they agree with the more canonically defined w_s.

Proof. Lemma 3.6 asserts that $\min_{j \leq n}\{sv_j(x) + j\} \geq \min_{j \leq n}\{sv_j^{\mathrm{naive}}(x) + j\}$, so it remains to prove the reverse inequality, which we do by induction on n. If $\min_{j \leq n}\{sv_j^{\mathrm{naive}}(x) + j\}$ is achieved by some $j < n$, then by the induction hypothesis,

$$\begin{aligned} \min_{j \leq n}\{sv_j^{\mathrm{naive}}(x) + j\} &= \min_{j \leq n - v_p(\pi)}\{sv_j^{\mathrm{naive}}(x) + j\} \\ &\geq \min_{j \leq n - v_p(\pi)}\{sv_j(x) + j\} \\ &\geq \min_{j \leq n}\{sv_j(x) + j\}. \end{aligned}$$

Suppose then that $\min_{j \leq n}\{sv_j^{\mathrm{naive}}(x) + j\}$ is achieved only for $j = n$. Put $x = \sum x_i u^i$; by definition, $v_n^{\mathrm{naive}}(x)$ is the smallest integer i with $v_p(x_i) \leq n$. In fact, we must have $v_p(x_i) = n$, or else we have $sv_j^{\mathrm{naive}}(x) + j < sv_n^{\mathrm{naive}}(x) + n$ for $j = v_p(x_i)$. Therefore $v_n(x_i u^i) = v_n^{\mathrm{naive}}(x_i u^i) = i$.

For $j < n$, $sv_j^{\mathrm{naive}}(x - x_i u^i) + j = sv_j^{\mathrm{naive}}(x) + j > si + n$. On the other hand, $v_n^{\mathrm{naive}}(x) = v_n^{\mathrm{naive}}(x_i u^i) = i$ and $v_n^{\mathrm{naive}}(x - x_i u^i) > i$. Thus for all $j \leq n$,

$$sv_j^{\mathrm{naive}}(x - x_i u^i) + j > si + n;$$

by Lemma 3.6, $sv_n(x - x_i u^i) + n > si + n$ and so $v_n(x - x_i u^i) > i = v_n(x_i u^i)$. Therefore $v_n(x) = v_n(x_i u^i) = i$, so

$$\min_{j \leq n}\{sv_j(x) + j\} \leq sv_n(x) + n = si + n = \min_{j \leq n}\{sv_j^{\mathrm{naive}}(x) + j\},$$

yielding the desired inequality. $\square$

Corollary 3.8. *We have $\Gamma^{\mathrm{alg}}_{\mathrm{con}} \cap \Gamma^{k((t))} = \Gamma^{k((t))}_{\mathrm{con}}$.*

We now define $\Gamma^K_{\mathrm{con}} = \Gamma^{\mathrm{alg}}_{\mathrm{con}} \cap \Gamma^K$, and Corollary 3.8 assures that this definition is consistent with our prior definition for $K = k((t))$. To show that $\Gamma^{\mathrm{alg}}_{\mathrm{con}} \cap \Gamma^K$ is "large" for any K, we need one more lemma, which will end up generalizing a standard fact about $\Gamma^{k((t))}_{\mathrm{con}}$.

Lemma 3.9. *For any valued field K, Γ^K_{con} is Henselian.*

Proof. By a lemma of Nagata [N, 43.2], it suffices to show that if $P(x) = x^d + a_1 x^{d-1} + \cdots + a_d$ is a polynomial over Γ^K_{con} such that $a_1 \not\equiv 0 \pmod{\pi}$ and $a_i \equiv 0 \pmod{\pi}$ for $i > 1$, then $P(x)$ has a root y in Γ^K_{con} such that $y \equiv -a_1 \pmod{\pi}$. By replacing $P(x)$ by $P(-x/a_1)$, we may reduce to the case $a_1 = -1$; by Hensel's lemma, P has a root y in Γ^K congruent to 1 modulo π, and $P'(y) \equiv dy^{d-1} - (d-1)y^{d-2} \equiv 1 \pmod{\pi}$.

Choose a constant $c > 0$ such that $v_n(a_i) \geq -cn$ for all n, and define the sequence $\{y_j\}_{j=0}^\infty$ by the Newton iteration, putting $y_0 = 1$ and $y_{j+1} = y_j - P(y_j)/P'(y_j)$. Then $\{y_j\}$ converges π-adically to y; we now show by induction on j that $v_n(y_j) \geq -cn$ for all n and all j. Namely, this is obvious for y_0, and given $v_n(y_j) \geq -cn$ for all n, it follows that $v_n(P(y_j)) \geq -cn$, $v_n(P'(y_j)) \geq -cn$, and $v_n(1/P'(y_j)) \geq -cn$ (the last because $v_0(P'(y_j)) = 0$). These together imply $v_n(y_{j+1}) \geq -cn$ for all n, completing the induction. We conclude that $y \in \Gamma^K_{\mathrm{con}}$ and Γ^K_{con} is Henselian, as desired. □

We can now prove the following.

Proposition 3.10. *For any valued field K, Γ^K_{con} has residue field K.*

Proof. We have already shown this for $K = k((t))$ by Corollary 3.8. If $K/k((t))$ is finite, then K uniquely determines L, M, N as in the definition of valued fields. Now $M = k_M((t))$ for some finite extension k_M of k, so Corollary 3.8 also implies that Γ^M_{con} has residue field M. Also, $N = M^{1/p^n}$ for some integer n, so for any $\overline{x} \in M$, we can find $y \in \Gamma^M_{\mathrm{con}}$ which lifts $\overline{x}^{p^n}$, and then $y^{\sigma^{-n}} \in \Gamma^N_{\mathrm{con}}$ lifts $\overline{x}$.

Choose a monic polynomial P over Γ^N_{con} lifting a monic separable polynomial $\overline{P}$ for which $K \cong N[x]/(\overline{P}(x))$ (again, possible by the primitive element theorem). The reduction is a separable polynomial, so by Hensel's lemma P has a root y in Γ^K, and $\Gamma^K \cong \Gamma^N[y]/(P(y))$. But since $\Gamma^{\mathrm{alg}}_{\mathrm{con}}$ is Henselian and P has coefficients in $\Gamma^{\mathrm{alg}}_{\mathrm{con}}$, $y \in \Gamma^{\mathrm{alg}}_{\mathrm{con}}$. Thus the residue field of Γ^K_{con} contains N and y, and hence is all of K.

This concludes the proof for K finite over $k((t))$. A general valued field K is the union of its finite valued subfields K_1, and Γ^K_{con} contains (but does not equal) the direct limit of the $\Gamma^{K_1}_{\mathrm{con}}$. Thus its residue field contains the union of the K_1, and hence is equal to K. □

If L/K is a finite extension of valued fields, then $\Gamma^L_{\mathrm{con}}/\Gamma^K_{\mathrm{con}}$ is an unramified extension of discrete valuation rings, and the corresponding residue field extension is L/K, so is finite. Thus Γ^L_{con} is integral over Γ^K_{con}.

3.3 Analytic rings: generalizing the Robba ring

In this section, we generalize the construction of the Robba ring. Besides the classical case where K is a finite extension of $k((t))$, we will be especially interested in the case

$K = k((t))^{\text{alg}}$, which will give a sort of "maximal unramified extension" of the standard Robba ring.

Proposition 3.11. *Suppose the valued field K is either*

(a) finite over $k((t))$ or

(b) perfect.

Then there exists $r > 0$ such that $\Gamma_r^K = \Gamma_r^{\text{alg}} \cap \Gamma^K$ has units congruent to every nonzero element of K.

Proof. For (a), let u be a lift to Γ_{con}^K of a uniformizer $\overline{u}$ of K, and choose $r > 0$ so that u is a unit in Γ_r^K. Let $\mathcal{O}'$ be the integral closure of $\mathcal{O}$ in Γ^K; its residue field is the integral closure k' of k in K.

For any $c_i \in \mathcal{O}'$, the series $1 + \sum_{i=1}^\infty c_i u^i$ converges with respect to w_r (hence levelwise) to a unit of Γ_r^K, because we can formally invert the series and the result also converges with respect to w_r. Any nonzero element of K can be written as a nonzero element of k' times a power of $\overline{u}$ times a series in $\overline{u}$ with leading term 1, thus can be lifted as an invertible element of $\mathcal{O}'$ times a power of u times a series of the form $1 + \sum_{i=1}^\infty c_i u^i$. The result is invertible in Γ_r^K, as desired.

For (b), we can choose any $r > 0$, since every Teichmüller lift belongs to Γ_r^K. ☐

Note that the conclusion of the proposition need not hold for other valued fields. For example, it fails for $K = k((t))^{\text{sep}}$ if $\sigma_0(u) = u^p$ for some $u \in \Gamma_{\text{con}}^{k((t))}$ lifting t: define a sequence $\{y_i\}_{i=1}^\infty$ of elements of K by setting y_i to be a root of $y_i^p - y_i = u^{-i}$. Then it can be shown that y_i has a lift in Γ_r^K only if $r < \frac{1}{i}(p/(p-1))^2$, so there is no way to choose r uniformly.

For the rest of this section, we assume that the hypotheses of Proposition 3.11 ace satisfied. Recall that for $0 < s \leq r$, we have defined the valuation w_s on $\Gamma_r^K[\frac{1}{p}]$ by

$$w_s(x) = \min_n \{n + s v_n(x)\},$$

the minimum taken as n runs over the value group of $\mathcal{O}$. We define a corresponding norm $|\cdot|_s$ by $|x|_s = p^{-w_s(x)}$.

While Γ_r^K is complete under $|\cdot|_r$, $\Gamma_r^K[\frac{1}{p}]$ is not, so we can attempt to complete it. In fact, we can define a Fréchet topology on $\Gamma_r^K[\frac{1}{p}]$ using the w_s for $0 < s \leq r$, and define $\Gamma_{\text{an},r}^K$ as the Fréchet completion of $\Gamma_r^K[\frac{1}{p}]$. That is, $\Gamma_{\text{an},r}^K$ consists of equivalence classes of sequences of elements of $\Gamma_r^K[\frac{1}{p}]$ which are simultaneously Cauchy for all of the norms $|\cdot|_s$.

Set $\Gamma_{\text{an,con}}^K = \cup_{r>0} \Gamma_{\text{an},r}^K$. Echoing a warning from the previous section, we note that $\Gamma_{\text{an,con}}^K$ admits an action of σ, but each $\Gamma_{\text{an},r}^K$ is mapped not into itself, but into $\Gamma_{\text{an},r/q}^K$. More precisely, we have $w_{r/q}(x^\sigma) = w_r(x)$ for all $x \in \Gamma_{\text{an},r}^K$.

In case $K = k((t))$, we defined another ring called $\Gamma_{\text{an,con}}^K$ in Section 2.3. Fortunately, these rings coincide: for r sufficiently small, by Corollary 3.8 we have $\Gamma_r^K = \Gamma_{r,\text{naive}}^K$ and so $\Gamma_{\text{an},r}^K = \Gamma_{\text{an},r,\text{naive}}^K$ by Proposition 2.2.

Since $\Gamma^K_{\mathrm{an,con}}$ is defined from Γ^K_{con} by a canonical completion process, it inherits as much functoriality as is possible given the restricted applicability of Proposition 3.11. For example, if L/K is a finite extension, then $\Gamma^L_{\mathrm{an,con}}$ is an integral extension of $\Gamma^K_{\mathrm{an,con}}$; in fact, one has a canonical identification of $\Gamma^L_{\mathrm{an,con}}$ with $\Gamma^L_{\mathrm{con}} \otimes_{\Gamma^K_{\mathrm{con}}} \Gamma^K_{\mathrm{an,con}}$, which in case L/K is Galois gives an action of $\mathrm{Gal}(L/K)$ on $\Gamma^L_{\mathrm{an,con}}$ with fixed ring $\Gamma^K_{\mathrm{an,con}}$. Likewise, if K is perfect, then the union $\cup_L \Gamma^L_{\mathrm{an},r}$ running over all finite subextensions L of K is dense in $\Gamma^K_{\mathrm{an},r}$ for each $r > 0$, so $\cup_L \Gamma^L_{\mathrm{an,con}}$ is dense in $\Gamma^K_{\mathrm{an,con}}$.

We can extend the functions v_n to $\Gamma^K_{\mathrm{an},r}$ by continuity: if $x_i \to x$ in the Fréchet topology, then $v_n(x_i)$ either stabilizes at some finite value or tends to $+\infty$ as $i \to \infty$, and we may put $v_n(x) = \lim_{i\to\infty} v_n(x_i)$. Likewise, we can extend the functions w_s to $\Gamma^K_{\mathrm{an},r}$ by continuity, and again one has the formula

$$w_s(x) = \min_n \{n + s v_n(x)\},$$

as n runs over the value group of $\mathcal{O}$. One also has

$$\lim_{n\to\pm\infty} (n + s v_n(x)) = \infty$$

for any $0 < s < r$. For $n \to -\infty$, this follows from the corresponding limiting statement for $s = r$. For $n \to \infty$, note that if the limit did not tend to infinity, x could not be written as a limit under $|\cdot|_s$ of elements of $\Gamma^K_r[\frac{1}{p}]$.

It is not so easy to prove anything about the ring $\Gamma^K_{\mathrm{an,con}}$ just from the above definition, since it is inconvenient to even write down elements of this ring. To this end, we isolate a special class of elements, which we call semiunits, and use them as building blocks to represent more general ring elements.

We define a *semiunit* of Γ^K_r (resp. of $\Gamma^K_{\mathrm{an},r}$) as an element u of Γ^K_r (resp. of $\Gamma^K_{\mathrm{an},r}$) which is either zero, or which satisfies the following conditions:

(a) $v_n(u) = \infty$ for $n < 0$;

(b) $v_0(u) < \infty$;

(c) $r v_n(u) + n > r v_0(u)$ for $n > 0$.

In particular, if $u \in \Gamma^K_r$, then u is a semiunit if either $u = 0$ or u is a unit in Γ^K_r, hence the terminology. In particular, under the condition of Proposition 3.11, every element of K lifts to a semiunit in Γ^K_r. Note that if u is a semiunit in $\Gamma^K_{\mathrm{an},r}$, it is also a semiunit in $\Gamma^K_{\mathrm{an},s}$ for any $0 < s < r$. Also beware that if $K/k((t))$ is infinite, a semiunit in $\Gamma^K_{\mathrm{an},r}$ need not belong to Γ^K_r even though $v_p(u) \geq 0$. (If R is the subring of $x \in \Gamma^K_{\mathrm{an},r}$ with $v_p(x) \geq 0$, then $R/\pi R$ is isomorphic to the completion of K with respect to v_K.)

If K is perfect, we define a *strong semiunit* of Γ^K_r (resp. of $\Gamma^K_{\mathrm{an},r}$) as an element u of Γ^K_r (resp. of $\Gamma^K_{\mathrm{an},r}$) which is either zero, or satisfies the following conditions:

(a) $v_n(u) = \infty$ for $n < 0$;

(b) $v_0(u) < \infty$;

(c) $v_n(u) = v_0(u)$ for $n > 0$.

Every Teichmüller lift is a strong semiunit, so every element of K lifts to a strong semiunit in Γ_r^K.

Let $\{u_i\}_{i=-\infty}^{\infty}$ be a doubly infinite sequence of semiunits in Γ_r^K (resp. in $\Gamma_{\mathrm{an},r}^K$). Then we say $\{u_i\}$ is a *semiunit decomposition* of x in Γ_r^K (resp. in $\Gamma_{\mathrm{an},r}^K$) if $w_r(u_i\pi^i) \leq w_r(u_j\pi^j)$ whenever $i > j$ and $u_i, u_j \neq 0$, and if $\sum_{i=-M}^{N} u_i\pi^i$ converges to x in the Fréchet topology as $M, N \to \infty$. We express this more succinctly by saying that $\sum u_i\pi^i$ is a semiunit decomposition of x. Analogously, if K is perfect and the u_i are strong semiunits, we say $\sum u_i\pi^i$ is a *strong semiunit decomposition* of x if $v_0(u_i) < v_0(u_j)$ whenever $i > j$ and $u_i, u_j \neq 0$, and if $\sum_{i=-M}^{N} u_i\pi^i$ converges to x in the Fréchet topology as $M, N \to \infty$.

If $\sum u_i\pi^i$ is a semiunit decomposition of $x \in \Gamma_{\mathrm{an},r}^K$, then for each i such that $u_i \neq 0$, we may set $n = iv_p(\pi)$ and obtain $rv_n(x) + n = rv_n(u_i\pi^i) + n$, that is, $v_n(x) = v_n(u_i\pi^i)$. Since $rv_n(x) + n \to \infty$ as $n \to \infty$ for any $x \in \Gamma_{\mathrm{an},r}^K$, we must then have $u_i = 0$ for i sufficiently large. There is no analogous phenomenon for strong semiunit decompositions, however: for each i such that $u_i \neq 0$, we set $n = iv_p(\pi)$ and obtain $v_n(x) = v_n(u_i\pi^i)$, but $v_n(x)$ may continue to decrease forever as $n \to \infty$, so the u_i need not eventually vanish.

Lemma 3.12. *Each element x of Γ_r^K admits a semiunit decomposition. If K is perfect, each element x of Γ_r^K admits a strong semiunit decomposition.*

Proof. Without loss of generality (by dividing by a suitable power of π), we may reduce to the case where $x \not\equiv 0 \pmod{\pi}$. We define a sequence of semiunits $\{y_i\}_{i=0}^{\infty}$ such that $x \equiv \sum_{i=0}^{j} y_i\pi^i \pmod{\pi^{j+1}}$, as follows. Let y_0 be a semiunit congruent to x modulo π. Given $y_0, \dots, y_j$, let y_{j+1} be a semiunit congruent to $(x - \sum_{i=0}^{j} y_i\pi^i)/\pi^{j+1}$ modulo π.

The sum $\sum_{i=0}^{\infty} y_i\pi^i$ now converges to x, but we do not have the necessary comparison between $w_r(y_i\pi^i)$ and $w_r(y_j\pi^j)$, so we must revise the decomposition. We say i is a *corner* if $w_r(y_i\pi^i) = \min_{j\leq i}\{w_r(y_j\pi^j)\}$. We now set $u_i = 0$ if i is not a corner; if i is a corner, let l be the next largest corner (or ∞ if there is none), and put $u_i = \sum_{j=i}^{l-1} y_j\pi^{j-i}$. By the definition of a corner, $w_r(y_j\pi^{j-i}) > w_r(y_i)$ for $i < j < l$, so u_i is a semiunit. Moreover, if i and j are corners and $i > j$, then $w_r(u_i\pi^i) = w_r(y_i\pi^i) \leq w_r(y_j\pi^j) = w_r(u_j\pi^j)$; and the sum $\sum_{i=0}^{\infty} u_i\pi^i$ is merely the sum $\sum_{i=0}^{\infty} y_i\pi^i$ with the terms regrouped, so it still converges to x. Thus $\sum_{i=0}^{\infty} u_i\pi^i$ is a semiunit decomposition of x.

If K is perfect, we perform the revision slightly differently. We say i is a corner if $v_0(y_i) < v_0(y_j)$ for all $j < i$. Again, we set $u_i = 0$ if i is not a corner, and if i is a corner and l is the next largest corner, we set $u_i = \sum_{j=i}^{l-i} y_j\pi^{j-i}$. Clearly u_i is a strong semiunit for each i, and the sum $\sum_{i=0}^{\infty} u_i\pi^i$ converges to x. If $i > j$ are corners, then $v_0(u_i) = v_0(y_i) < v_0(y_j) = v_0(u_j)$. Thus $\sum_{i=0}^{\infty} u_i\pi^i$ is a strong semiunit decomposition of x. □

Proposition 3.13. *Every element of $\Gamma_{\mathrm{an},r}^K$ admits a semiunit decomposition.*

Proof. For $x \in \Gamma_{\mathrm{an},r}^K$, let $\sum_{l=0}^{\infty} x_l$ be a series of elements of $\Gamma_r^K[\frac{1}{p}]$ that converges under $|\cdot|_r$ to x, such that $w_r(x_l) < w_r(x_{l+1})$. (For example, choose x_0 such that $w_r(x - x_0) > w_r(x)$, then choose x_1 such that $w_r(x - x_0 - x_1) > w_r(x - x_0)$, and so forth.)

For $l = 0, 1, \dots$ and $i \in \mathbb{Z}$, we define elements y_{il} of $\Gamma_r^K[\frac{1}{p}]$ recursively in l, such that for any l, only finitely many of the y_{il} are nonzero, as follows. Apply Lemma 3.12 (after multiplying by a suitable power of π) to produce a semiunit decomposition of $x_0 + \cdots + x_l - \sum_{j<l} \sum_i y_{ij}\pi^i$. For each of the finitely many terms $u_i\pi^i$ of this decomposition with $u_i \neq 0$ and $w_r(u_i\pi^i) < w_r(x_{l+1})$, put $y_{il} = u_i$; for all other i, put $y_{il} = 0$. Then

$$w_r\left(x_0 + \cdots + x_l - \sum_{j\leq l}\sum_i y_{ij}\pi^i\right) \geq w_r(x_{l+1}).$$

In particular, the doubly infinite sum $\sum_l \sum_i y_{il}\pi^i$ converges under $|\cdot|_r$ to x. If we set $z_i = \sum_l y_{il}$, the series $\sum_i z_i\pi^i$ converges under $|\cdot|_r$ to x.

Note that $w_r(x_l) \leq w_r(y_{il}\pi^i) < w_r(x_{l+1})$ whenever $y_{il} \neq 0$. Thus for any fixed i, the values of $w_r(y_{il}\pi^i)$, taken over all l such that $y_{il} \neq 0$, form a strictly increasing sequence. If j is the first such index, we then have $w_r(y_{ij}\pi^i) < w_r(\sum_{l>j} y_{il}\pi^i)$, and so z_i is a semiunit.

Define u_i to be zero if $w_r(z_i\pi^i) > w_r(z_j\pi^j)$ for some $j < i$; otherwise, let l be the smallest integer greater than i such that $w_r(z_l\pi^l) \leq w_r(z_i\pi^i)$ (or ∞ if none exists), and put $u_i = \sum_{j=i}^{l-1} z_j\pi^{j-i}$. Then the series $\sum_i u_i\pi^i$ also converges under $|\cdot|_r$ to x, and if $u_i \neq 0$, then $w_r(u_i\pi^i) = w_r(z_i\pi^i)$. It follows that $w_r(u_i\pi^i) \leq w_r(u_j\pi^j)$ whenever $i > j$ and $u_i, u_j \neq 0$. This in turn implies that if $u_i \neq 0$ and $n = v_p(\pi^i)$, then $v_n(u_i\pi^i) = v_n(x)$.

We finally check that $\sum_i u_i\pi^i$ converges under $|\cdot|_s$ for $0 < s < r$. The fact that $sv_n(x) + n \to \infty$ as $n \to \pm\infty$ implies that $sv_{v_p(\pi^i)}(u_i\pi^i) + v_p(\pi^i) \to \infty$ as $i \to \pm\infty$. Since u_i is a semiunit, $w_s(u_i\pi^i) = sv_{v_p(\pi^i)}(u_i\pi^i) + v_p(\pi^i)$, so $w_s(u_i\pi^i) \to \infty$ as $i \to \pm\infty$. Thus the sum $\sum_i u_i\pi^i$ converges under $|\cdot|_s$ for $0 < s < r$, and the limit must equal x because the sum converges to x under $|\cdot|_r$. Therefore $\sum_i u_i\pi^i$ is a semiunit decomposition, as desired. □

Proposition 3.14. *If K is perfect, every element of $\Gamma_{\mathrm{an},r}^K$ admits a strong semiunit decomposition.*

Proof. As in the previous proof, for $x \in \Gamma_{\mathrm{an},r}^K$, let $\sum_{l=0}^\infty x_l$ be a series of elements of $\Gamma_r^K[\frac{1}{p}]$ that converges under $|\cdot|_r$ to x, such that $w_r(x_l) < w_r(x_{l+1})$.

For $l = 0, 1, \dots$ and $i \in \mathbb{Z}$, we define elements y_{il} of $\Gamma_r^K[\frac{1}{p}]$ recursively in l, such that for any l, only finitely many of the y_{il} are nonzero, as follows. Apply Lemma 3.12 to produce a strong semiunit decomposition of $x_0 + \cdots + x_l - \sum_{j<l}\sum_i y_{ij}\pi^i$. For each of the finitely many terms $u_i\pi^i$ of this decomposition with $u_i \neq 0$ and $w_r(u_i\pi^i) < w_r(x_{l+1})$, put $y_{il} = u_i$; for all other i, put $y_{il} = 0$. Then

$$w_r\left(x_0 + \cdots + x_l - \sum_{j\leq l}\sum_i y_{ij}\pi^i\right) \geq w_r(x_{l+1}).$$

In particular, the doubly infinite sum $\sum_l \sum_i y_{il}\pi^i$ converges under $|\cdot|_r$ to x. If we set $z_i = \sum_l y_{il}$, and the series $\sum_i z_i\pi^i$ converges under $|\cdot|_r$ to x.

Note that $w_r(x_l) \leq w_r(y_{il}\pi^i) < w_r(x_{l+1})$ whenever $y_{il} \neq 0$. Thus for any fixed i, the values of $v_0(y_{il})$, taken over all l such that $y_{il} \neq 0$, form a strictly increasing sequence. If j is the first such index, we then have $v_0(y_{ij}) < v_0(\sum_{l>j} y_{il})$, and so z_i is a strong semiunit.

Define u_i to be zero if $v_0(z_i) \geq v_0(z_j)$ for some $j < i$; otherwise, let l be the smallest integer such that $v_0(z_l) < v_0(z_i)$ (or ∞ if none exists), and put $u_i = \sum_{j=i}^{l-1} z_j \pi^{j-i}$. Then the series $\sum_i u_i \pi^i$ also converges under $|\cdot|_r$ to x, and if $u_i \neq 0$, then $v_0(u_i) = v_0(z_i)$. It follows that $v_0(u_i) < v_0(u_j)$ whenever $i > j$ and $u_i, u_j \neq 0$. This in turn implies that if $u_i \neq 0$ and $n = v_p(\pi^i)$, then $v_n(u_i \pi^i) = v_n(x)$.

We finally check that $\sum_i u_i \pi^i$ converges under $|\cdot|_s$ for $0 < s < r$, by the same argument as in the previous proof. Namely, the fact that $s v_n(x) + n \to \infty$ as $n \to \pm\infty$ implies that $s v_{v_p(\pi^i)}(u_i \pi^i) + v_p(\pi^i) \to \infty$ as $i \to \pm\infty$. Since u_i is a strong semiunit, $w_s(u_i \pi^i) = s v_{v_p(\pi^i)}(u_i \pi^i) + v_p(\pi^i)$, so $w_s(u_i \pi^i) \to \infty$ as $i \to \pm\infty$. Thus the sum $\sum_i u_i \pi^i$ converges under $|\cdot|_s$ for $0 < s < r$, and the limit must equal x because the sum converges to x under $|\cdot|_r$. Therefore $\sum_i u_i \pi^i$ is a strong semiunit decomposition, as desired. ☐

Although (strong) semiunit decompositions are not unique, in a certain sense the "leading terms" are unique. To make sense of this remark, we first need a "leading coefficient map" for K.

Lemma 3.15. *For K a valued field, there exists a homomorphism $\lambda : K^* \to (k^{\mathrm{alg}})^*$ such that $\lambda(c) = c$ for all $c \in k^{\mathrm{alg}} \cap K$ and $\lambda(x) = 1$ if $v_K(x-1) > 0$.*

For instance, if $K = k((t))$, we could take $\lambda(x)$ to be the leading coefficient of x.

Proof. There is no loss of generality in enlarging K, so we may assume $K = k((t))^{\mathrm{alg}}$. Define $t_0 = t$, and for $i > 0$, let t_i be an i-th root of t_{i-1}. With this choice, for any $d \in \mathbb{Q}$ we can define t^d as $t_i^{i!d}$ for any $i \geq d$; the expression does not depend on i.

Now for each $x \in K^*$, there exists a unique $c \in k^{\mathrm{alg}}$ such that

$$v_K\left(\frac{x}{c t^{v_K(x)}} - 1\right) > 0;$$

set $\lambda(x) = c$. ☐

Choose a map λ as in Lemma 3.15; we define the leading terms map $L_r : \Gamma^K_{\mathrm{an},r} \to \cup_{n=1}^{\infty} k^{\mathrm{alg}}[t^{1/n}, t^{-1/n}]$ as follows. For $x \in \Gamma^K_{\mathrm{an},r}$ nonzero, find a finite sum $y = \sum_j u_j \pi^j$ such that each u_j is a semiunit, $w_r(u_j \pi^j) = w_r(x)$ for all j such that $u_j \neq 0$, and $w_r(x-y) > w_r(x)$. Then put $L_r(x) = \sum_j \lambda(\overline{u_j}) t^{v_0(u_j)}$; this definition does not depend on the choice of y. Moreover, the leading terms map is multiplicative, that is, $L_r(xy) = L_r(x) L_r(y)$.

We define the *upper degree* and *lower degree* of a nonzero element of $\cup_{n=1}^{\infty} k^{\mathrm{alg}}[t^{1/n}, t^{-1/n}]$ as the largest and smallest powers of t, respectively, occurring in the element; we define the *length* of an element as the upper degree minus the lower degree. We extend all of these definitions to $\Gamma^K_{\mathrm{an},r}$ through the map L_r.

Warning: if K is not finite over $k((t))$, then the subring of $x \in \Gamma^K_{\mathrm{an,con}}$ with $v_n(x) = \infty$ for $n < 0$ is a complete discrete valuation ring containing Γ^K_{con}, but it is actually much bigger than Γ^K_{con}. In fact, its residue field is the completion of K with respect to the valuation v_K.

As noted earlier, a theorem of Lazard asserts that $\Gamma^K_{\mathrm{an,con}}$ is a Bézout ring (every finitely generated ideal is principal) for $K = k((t))$; the same is true for K a finite extension of $k((t))$, since $K \cong k'((t'))$ for some uniformizer t' and some field k'. We will generalize the Bézout property to $\Gamma^K_{\mathrm{an,con}}$ for $K/k((t))$ infinite in Section 3.6; for now, we deduce from Lemma 2.6 the following descent lemma for σ-modules. (The condition on G-stable ideals is satisfied because $G = \mathrm{Gal}(L/K)$ here is finite.)

Corollary 3.16. *Let L/K be a finite Galois extension of valued fields finite over $k((t))$. Let M be a σ-module over $\Gamma^K_{\mathrm{an,con}}$ and N a saturated σ-submodule of $M \otimes_{\Gamma^K_{\mathrm{an,con}}} \Gamma^L_{\mathrm{an,con}}$ stable under $\mathrm{Gal}(L/K)$. Then N is equal to $P \otimes_{\Gamma^K_{\mathrm{an,con}}} \Gamma^L_{\mathrm{an,con}}$ for some saturated σ-submodule P of M.*

3.4 Some σ-equations

We record here the behavior of some simple equations involving σ. For starters, we have the following variant of Hensel's lemma.

Proposition 3.17. *Let R be a complete discrete valuation ring, unramified over $\mathcal{O}$, with separably closed residue field, and let σ be a q-power Frobenius lift. For $c_0, \dots, c_n \in R$ with c_0 not divisible by π and $x \in R$, define $f(x) = c_0 x + c_1 x^\sigma + \cdots + c_n x^{\sigma^n}$. Then for any $x, y \in R$ for which $f(x) \equiv y \pmod{\pi}$, there exists $z \in R$ congruent to x modulo π for which $f(z) = y$. Moreover, if R has algebraically closed residue field, then the same holds if any of $c_0, \dots, c_n$ is not divisible by π.*

Proof. Define a sequence $\{z_l\}_{l=1}^\infty$ of elements of R such that $z_1 = x$, $z_{l+1} \equiv z_l \pmod{\pi^l}$ and $f(z_l) \equiv y \pmod{\pi^l}$; then the limit z of the z_l will have the desired property. Given z_l, put $a_l = (y - f(z_l))/\pi^l$, and choose $b_l \in R$ such that

$$c_0 b_l + c_1 b_l^q (\pi^\sigma/\pi)^l + \cdots + c_n b_l^{q^n} (\pi^{\sigma^n}/\pi)^l \equiv a_l \pmod{\pi};$$

this is possible because either R has algebraically closed residue field, or $c_0 \neq 0$ and the polynomial at left must be separable. Put $z_{l+1} = z_l + \pi^l b_l$; then $f(z_{l+1}) \equiv f(z_l) + f(\pi^l b_l) \equiv y \pmod{\pi^{l+1}}$, as desired. ☐

We next consider similar equations over some other rings. The following result will be vastly generalized by Proposition 5.11 later.

Proposition 3.18. *Suppose $x \in \Gamma^{\mathrm{alg}}_{\mathrm{con}}$ (resp. $x \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ with $v_n(x) = \infty$ for $n < 0$) is not congruent to 0 modulo π. Then there exists a nonzero $y \in \Gamma^{\mathrm{alg}}_{\mathrm{con}}$ (resp. $y \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ with $v_n(y) = \infty$ for $n < 0$) such that $y^\sigma = xy$.*

Proof. Put $R = \Gamma^{\mathrm{alg}}_{\mathrm{con}}$ (resp. let R be the subring of $x \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ with $v_n(x) = \infty$ for $n < 0$) and let S be the completion of R. By Proposition 3.17, we can find nonzero $y \in S$ such that $y^\sigma = xy$; we need to show that $y \in R$. Choose $r > 0$ and $c \in \mathbb{R}$ such that $r v_n(x) + n \geq c$ for all n. We then show that $r(q-1)v_n(y) + n \geq c$ by induction on n. We have

$$q v_n(y) = v_n(y^\sigma) \geq \min_{m \leq n} \{v_m(x) + v_{n-m}(y)\}.$$

If the minimum is achieved for $m = 0$ (which includes the base case $n = 0$), then $(q - 1)v_n(y) \geq v_0(x)$, so $r(q-1)v_n(y) + n \geq r(q-1)v_0(x) + n \geq c$. If the minimum is achieved for some $m > 0$, then by the induction hypothesis

$$\begin{aligned} r(q-1)v_n(y) &\geq \frac{r(q-1)}{q} v_m(x) + \frac{r(q-1)}{q} v_{n-m}(y) \\ &\geq \frac{(q-1)(c-m)}{q} + \frac{(c-n+m)}{q} \\ &\geq \frac{(q-1)(c-n)}{q} + \frac{(c-n)}{q} \geq c - n, \end{aligned}$$

so $r(q-1)v_n(y) + n \geq c$. Thus the induction goes through, and demonstrates that $y \in R$, as desired. □

Finally, we consider a class of equations involving the analytic rings. We suppress K from all superscripts for convenience, writing Γ_{con} for Γ^K_{con} and so forth.

Proposition 3.19. *Let K be a valued field (satisfying the condition of Proposition 3.11 in case $\Gamma_{\mathrm{an,con}}$ is referenced).*

(a) Assume K is separably closed (resp. algebraically closed). For $\lambda \in \mathcal{O}$ a unit and $x \in \Gamma_{\mathrm{con}}$ (resp. $x \in \Gamma_{\mathrm{an,con}}$), there exists $y \in \Gamma_{\mathrm{con}}$ (resp. $y \in \Gamma_{\mathrm{an,con}}$) such that $y^\sigma - \lambda y = x$. Moreover, if $x \in \Gamma_{\mathrm{con}}[\frac{1}{p}]$, then any such y belongs to $\Gamma_{\mathrm{con}}[\frac{1}{p}]$.

(b) Assume K is perfect. For $\lambda \in \mathcal{O}$ not a unit and $x \in \Gamma_{\mathrm{con}}$ (resp. $x \in \Gamma_{\mathrm{an,con}}$), there exists $y \in \Gamma_{\mathrm{con}}$ (resp. $y \in \Gamma_{\mathrm{an,con}}$) such that $y^\sigma - \lambda y = x$. Moreover, we can take y nonzero in $\Gamma_{\mathrm{an,con}}$ even if $x = 0$.

(c) For $\lambda \in \mathcal{O}$ not a unit and $x \in \Gamma_{\mathrm{an,con}}$, there is at most one $y \in \Gamma_{\mathrm{an,con}}$ such that $\lambda y^\sigma - y = x$, and if $x \in \Gamma_{\mathrm{con}}$, then $y \in \Gamma_{\mathrm{con}}$ as well.

(d) For $\lambda \in \mathcal{O}$ not a unit and $x \in \Gamma_{\mathrm{an,con}}$ such that $v_n(x) \geq 0$ for all n, there exists $y \in \Gamma_{\mathrm{an,con}}$ such that $\lambda y^\sigma - y = x$.

Proof. (a) If $x \in \Gamma_{\mathrm{con}}$, then Proposition 3.17 implies that there exists $y \in \Gamma$ such that $y^\sigma - \lambda y = x$. To see that in fact $y \in \Gamma_{\mathrm{con}}$, note that if $v_n(y) \leq 0$, the fact that

$$q v_n(y) = v_n(y^\sigma) = v_n(\lambda y + x) \geq \min\{v_n(x), v_n(y)\}$$

implies that $q v_n(y) \geq v_n(x)$; while if $v_n(y) > 0$, the fact that

$$v_n(y) = v_n(\lambda y) = v_n(y^\sigma - x) \geq \min\{q v_n(y), v_n(x)\}$$

implies that $v_n(y) \geq v_n(x)$, which also implies $q v_n(y) \geq v_n(x)$. Hence $y \in \Gamma_{\mathrm{con}}$ and $w_{qr}(y) \geq w_r(x)$.

For $x \in \Gamma_{\mathrm{an,con}}$ (with K algebraically closed), choose $r > 0$ such that $x \in \Gamma_{\mathrm{an},r}$, and let $x = \sum_{i=-\infty}^{\infty} u_i \pi^i$ be a strong semiunit decomposition. As above, there exists $y_i \in \Gamma_{\mathrm{an},qr}$ with $y_i^\sigma - \lambda(\pi/\pi^\sigma)^i y_i = u_i(\pi/\pi^\sigma)^i$ such that $v_n(y_i) = \infty$ for $n < 0$ and $w_{qr}(y_i) \geq w_r(u_i)$. This implies that $\sum_{i=-\infty}^{\infty} y_i \pi^i$ converges with respect to $|\cdot|_s$ for $0 < s \leq r$; let y be its Fréchet limit. Then

$$\begin{aligned} y^\sigma - \lambda y &= \sum_i y_i^\sigma (\pi^i)^\sigma - \lambda y_i \pi^i \\ &= \sum_i \lambda y_i \pi^i + u_i \pi^i - \lambda y_i \pi^i \\ &= \sum_i u_i \pi^i = x, \end{aligned}$$

so y is the desired solution.

To verify the last assertion, suppose $x \in \Gamma_{\mathrm{con}}[\frac{1}{p}]$ and $y \in \Gamma_{\mathrm{an,con}}$ satisfy $y^\sigma - \lambda y = x$. By what we have shown above, there also exists $z \in \Gamma_{\mathrm{con}}[\frac{1}{p}]$ such that $z^\sigma - \lambda z = x$, so $(y-z)^\sigma = \lambda(y-z)$. This equation yields $q v_n(y-z) = v_n(y-z)$ for all n, so $v_n(y-z) = 0$ or ∞ for all n. We cannot have $v_n(y-z) = 0$ for all n, so there is a smallest such n; we may assume $n = 0$ without loss of generality. Let $\mathcal{O}'$ be the completed integral closure of $\mathcal{O}$ in Γ_{con}. Then every solution w of $w^\sigma = \lambda w$ in $\Gamma_{\mathrm{an,con}}$ with $v_n(w) = \infty$ for $n < 0$ is congruent to some element of $\mathcal{O}'$ modulo π. In particular, we can find $c_0, c_1, \dots \in \mathcal{O}'$ such that $\sum_{j=0}^{l} c_j \pi^j \equiv y - z \pmod{\pi^{l+1}}$, since once $c_0, \dots, c_l$ have been computed, we can take $w = (y-z)\pi^{-l-1} - \sum_{j=0}^{l} c_j \pi^{j-l-1}$, and there must be some $c_{l+1} \in \mathcal{O}'$ congruent to w modulo π. Thus $y - z \in \mathcal{O}' \subseteq \Gamma_{\mathrm{con}}[\frac{1}{p}]$, so $y \in \Gamma_{\mathrm{con}}[\frac{1}{p}]$.

(b) If $x \in \Gamma_{\mathrm{con}}$, then the series

$$\sum_{i=0}^{\infty} \lambda^{\sigma^{-1}} \cdots \lambda^{\sigma^{-i}} x^{\sigma^{-i-1}}$$

converges π-adically to an element $y \in \Gamma$ satisfying

$$\begin{aligned} y^\sigma - \lambda y &= \sum_{i=0}^{\infty} \lambda \lambda^{\sigma^{-1}} \cdots \lambda^{\sigma^{-i+1}} x^{\sigma^{-i}} - \sum_{i=0}^{\infty} \lambda \lambda^{\sigma^{-1}} \cdots \lambda^{\sigma^{-i}} x^{\sigma^{-i-1}} \\ &= \sum_{i=0}^{\infty} \lambda \lambda^{\sigma^{-1}} \cdots \lambda^{\sigma^{-i+1}} x^{\sigma^{-i}} - \sum_{i=1}^{\infty} \lambda \lambda^{\sigma^{-1}} \cdots \lambda^{\sigma^{-i+1}} x^{\sigma^{-i}} \\ &= x. \end{aligned}$$

To see that in fact $y \in \Gamma_{\mathrm{con}}$, choose $r > 0$ and $c \leq 0$ such that $w_r(x) \geq c$, that is, $r v_n(x) + n \geq c$ for all $n \geq 0$. If $v_n(x) \leq 0$, then $r v_n(x^{\sigma^{-i}}) + n = (r/q^i) v_n(x) + n \geq r v_n(x) + n \geq c$; if $v_n(x) \geq 0$, then $r v_n(x^{\sigma^{-i}}) + n \geq 0 \geq c$. In any case, we have $w_r(x^{\sigma^{-i}}) \geq c$ for all i. Since $w_r(\lambda^{\sigma^{-i}}) = w_r(\lambda) > 0$ for all i, we conclude that the series defining y converges under $|\cdot|_r$, and so its limit y in Γ must actually lie in Γ_{con}.

Suppose now that $x \in \Gamma_{\mathrm{an,con}}$; by Proposition 3.14, there exists a strong semiunit decomposition $x = \sum_n \pi^n u_n$ of x. Let N be the largest value of n for which $v_0(u_n) \geq 0$, and put

$$x_- = \sum_{n=-\infty}^{N} \pi^n u_n, \qquad x_+ = \sum_{n=N+1}^{\infty} \pi^n u_n.$$

As above, we can construct $y_+ \in \Gamma_{\mathrm{con}}[\frac{1}{p}]$ so that $y_+^\sigma - \lambda y_+ = x_+$. As for x_-, let m be the greatest integer less than or equal to N for which $u_m \neq 0$. For any fixed r, $w_r(x_-^{\sigma^i}) = w_r((u_m \pi^m)^{\sigma^i})$ for i sufficiently large. The series

$$-\sum_{i=0}^{\infty} (\lambda \lambda^\sigma \cdots \lambda^{\sigma^i})^{-1} x_-^{\sigma^i}$$

then converges under $|\cdot|_r$, since

$$\begin{aligned} w_r((\lambda \lambda^\sigma \cdots \lambda^{\sigma^i})^{-1} x_-^{\sigma^i}) &= -(i+1) w_r(\lambda) + w_r(x_-^{\sigma^i}) \\ &= -(i+1) w_r(\lambda) + r q^i v_0(u_m) + m v_p(\pi) \end{aligned}$$

tends to infinity with i. Since this holds for every r, the series converges in $\Gamma_{\mathrm{an,con}}$ to a limit y_-, which satisfies

$$\begin{aligned} y_-^\sigma - \lambda y_- &= -\sum_{i=0}^{\infty} (\lambda^\sigma \cdots \lambda^{\sigma^{i+1}})^{-1} x_-^{\sigma^{i+1}} + \sum_{i=0}^{\infty} (\lambda^\sigma \cdots \lambda^{\sigma^i})^{-1} x_-^{\sigma^i} \\ &= -\sum_{i=1}^{\infty} (\lambda^\sigma \cdots \lambda^{\sigma^i})^{-1} x_-^{\sigma^i} + \sum_{i=0}^{\infty} (\lambda^\sigma \cdots \lambda^{\sigma^i})^{-1} x_-^{\sigma^i} \\ &= x_-. \end{aligned}$$

We conclude that $y = y_+ + y_-$ satisfies $y^\sigma - \lambda y = x$.

To prove the final assertion, let u be any strong semiunit with $v_0(u) > 0$, and set

$$y = \sum_{i=0}^{\infty} \lambda^{\sigma^{-1}} \cdots \lambda^{\sigma^{-i}} u^{\sigma^{-i-1}} + \sum_{i=0}^{\infty} (\lambda \lambda^\sigma \cdots \lambda^{\sigma^i})^{-1} u^{\sigma^i};$$

then the above arguments show that both series converge and $y^\sigma - \lambda y = u - u = 0$.

(c) We prove the second assertion first. Namely, assume $x \in \Gamma_{\mathrm{con}}$ and $y \in \Gamma_{\mathrm{an,con}}$ satisfy $\lambda y^\sigma - y = x$; we show that $y \in \Gamma_{\mathrm{con}}$. First suppose $0 < v_n(y) < \infty$ for some $n < 0$. Then

$$v_n(y) = v_n(y + x) = v_n(\lambda y^\sigma) \geq v_n(y^\sigma) = q v_n(y),$$

contradiction. Thus $v_n(y)$ is either nonpositive or ∞ for all $n < 0$. We cannot have $v_n(y) < 0$ for all n, since for some $r > 0$ we have $r v_n(y) + n \to \infty$ as $n \to -\infty$. Thus

$v_n(y) = \infty$ for some y. (Beware: this is not enough *a priori* to imply that $y \in \Gamma_{\mathrm{con}}[\frac{1}{p}]$ if K is infinite over $k((t))$.) Choose n minimal such that $v_n(y) < \infty$. If $n < 0$, then $v_n(y) = v_n(y + x) = v_n(\lambda y^\sigma) = \infty$, contradiction. Thus $n \geq 0$. We can now show that y is congruent modulo π^i to an element of Γ_{con}, by induction on i. The base case $i = 0$ is vacuous; given $y \equiv y_i \pmod{\pi^i}$ for $y_i \in \Gamma_{\mathrm{con}}$, we have

$$y = -x + \lambda y^\sigma \equiv -x + \lambda y_i^\sigma \pmod{\pi^{i+1}}.$$

Thus the induction follows. Since y is the π-adic limit of elements of Γ_{con}, we conclude $y \in \Gamma_{\mathrm{con}}$.

For the first assertion, suppose $x \in \Gamma_{\mathrm{an,con}}$ and $y_1, y_2 \in \Gamma_{\mathrm{an,con}}$ satisfy $\lambda y_i^\sigma - y_i = x$ for $i = 1, 2$. Then $\lambda(y_1 - y_2)^\sigma - (y_1 - y_2) = 0$; by the previous paragraph, this implies $y_1 - y_2 \in \Gamma_{\mathrm{con}}$. But then $v_p(y_1 - y_2) = v_p(\lambda) + v_p((y_1 - y_2)^\sigma)$, a contradiction unless $y_1 - y_2 = 0$.

(d) Since $v_n(x) \geq 0$ for all n, we have $v_n(x^{\sigma^i}) = q^i v_n(x) \geq v_n(x)$ for all nonnegative integers i. Thus $w_s(x^{\sigma^i}) \geq w_s(x)$ for all s, so the series

$$y = -\sum_{i=0}^{\infty} \lambda \lambda^\sigma \cdots \lambda^{\sigma^{i-1}} x^{\sigma^i}$$

converges with respect to each of the norms $|\cdot|_s$, and

$$\begin{aligned}
\lambda y^\sigma - y &= -\sum_{i=0}^{\infty} \lambda \lambda^\sigma \cdots \lambda^{\sigma^i} x^{\sigma^{i+1}} + \sum_{i=0}^{\infty} \lambda \lambda^\sigma \cdots \lambda^{\sigma^{i-1}} x^{\sigma^i} \\
&= -\sum_{i=1}^{\infty} \lambda \lambda^\sigma \cdots \lambda^{\sigma^{i-1}} x^{\sigma^i} + \sum_{i=0}^{\infty} \lambda \lambda^\sigma \cdots \lambda^{\sigma^{i-1}} x^{\sigma^i} \\
&= x,
\end{aligned}$$

so y is the desired solution.

□

3.5 Factorizations over analytic rings

We assume that the valued field K satisfies the conditions of Proposition 3.11, so that the ring $\Gamma_{\mathrm{an,con}} = \Gamma^K_{\mathrm{an,con}}$ is defined. As noted earlier, $\Gamma_{\mathrm{an,con}}$ is not Noetherian even for $K = k((t))$, but in this case Lazard [L] proved that $\Gamma_{\mathrm{an,con}}$ is a Bézout ring, that is, a ring in which every finitely generated ideal is principal. In this section and the next, we generalize Lazard's result as follows.

Theorem 3.20. *Suppose the conclusion of Proposition 3.11 is satisfied for the valued field K and the positive number r. Then every finitely generated ideal in $\Gamma_{\mathrm{an},r} = \Gamma^K_{\mathrm{an},r}$ is principal. In particular, every finitely generated ideal in $\Gamma_{\mathrm{an,con}}$ is principal.*

Our approach resembles that of Lazard, with "pure elements" standing in for the divisors in his theory. The approach requires a number of auxiliary results on factorizations of elements of $\Gamma_{\mathrm{an,con}}$; for the most part (specifically, excepting Section 6.1), only Theorem 3.20 will be used in the sequel, not the auxiliary results.

For $x \in \Gamma_{\mathrm{an},r}$ nonzero, define the *Newton polygon* of x as the lower convex hull of the set of points $(v_n(x), n)$, minus any segments of slopes less than $-r$ on the left end and/or any segments of nonnegative slope on the right end of the polygon; see Figure 1 for an example. Define the *slopes* of x as the negatives of the slopes of the Newton polygon of x. (The negation is to ensure that the slopes of x are positive.) Also define the *multiplicity* of a slope $s \in (0, r]$ of x as the positive difference in y-coordinates between the endpoints of the segment of the Newton polygon of slope $-s$, or 0 if there is no such segment. If x has only one slope s, we say x is *pure* (of slope s). (Beware: this notion of slope differs from the slope of an eigenvector of a σ-module introduced in Section 2.5, and the Newton polygon here does not correspond to either the generic or special Newton polygons we define later.)

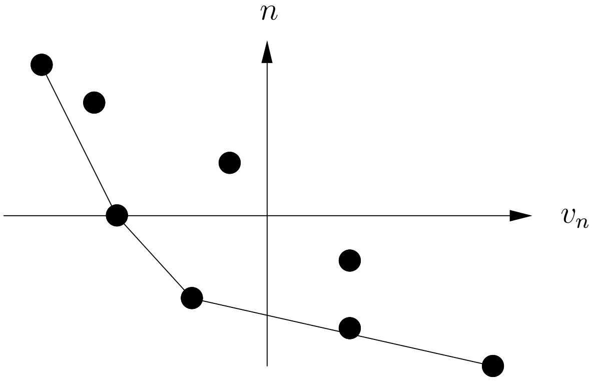

Figure 1: An example of a Newton polygon

Lemma 3.21. *The multiplicity of s as a slope of x is equal to s times the length (upper degree minus lower degree) of $L_s(x)$, where L_s is the leading terms map in $\Gamma_{\mathrm{an},s}$.*

Proof. Let $\sum_i u_i \pi^i$ be a semiunit decomposition of x. Let S be the set of l which achieve $\min_l\{w_s(u_l \pi^l)\}$, and let i and j be the smallest and largest elements of S; then $L_s(x) = \sum_{l \in S} \lambda(\overline{u_l}) t^{v_0(u_l)}$ and the length of $L_s(x)$ is equal to $v_0(u_i) - v_0(u_j)$.

We now show that the endpoints of the segment of the Newton polygon of x of slope $-s$ are $(v_0(u_i), v_p(\pi^i))$ and $(v_0(u_j), v_p(\pi^j))$. First of all, for $n = v_p(\pi^i)$, we have $sv_n(x) + n = sv_0(u_i) + iv_p(\pi) = w_s(u_i\pi^i)$; likewise for $n = v_p(\pi^j)$. Next, we note that $w_s(x) \geq \min_l\{w_s(u_l\pi^l)\} = w_s(u_i\pi^i)$. Thus for any n, $sv_n(x) + n \geq w_s(u_i\pi^i)$; this means that the line through $(v_0(u_i), v_p(\pi^i))$ and $(v_0(u_j), v_p(\pi^j))$ is a lower supporting line for the set of points

$(v_n(x), n)$. Finally, note that for $n < v_p(\pi^i)$,

$$\begin{aligned} sv_n(x) + n &\geq \min_{l<i}\{sv_n(\pi^l u_l) + n\} \\ &\geq \min_{l<i}\{w_s(\pi^l u_l)\} \\ &> w_s(x); \end{aligned}$$

while for $n > v_p(\pi^j)$,

$$\begin{aligned} sv_n(x) + n &\geq \min\{\min_{l\in[i,j]}\{sv_n(\pi^l u_l) + n\}, \min_{l\notin[i,j]}\{sv_n(\pi^l u_l) + n\}\} \\ &\geq \min\{\min_{l\in[i,j]}\{sv_n(\pi^l u_l) + n\}, \min_{l\notin[i,j]}\{w_s(\pi^l u_l)\}\}. \end{aligned}$$

For $l \in [i,j]$, $n > v_p(\pi^l)$ and u_l is a semiunit, so $sv_n(\pi^l u_l) + n > w_s(\pi^l u_l) \geq w_s(x)$; for $l \notin [i,j]$, $w_s(\pi^l u_l) > w_s(x)$ by the choice of i and j. Putting the inequalities together, we again conclude $sv_n(x) + n > w_s(x)$.

Therefore the endpoints of the segment of the Newton polygon of x of slope $-s$ are $(v_0(u_i), v_p(\pi^i))$ and $(v_0(u_j), v_p(\pi^j))$. Thus the multiplicity of s as a slope of x is $v_p(\pi^j) - v_p(\pi^i) = s(v_0(u_i) - v_0(u_j))$, which is indeed s times the length of $L_s(x)$, as claimed. □

Corollary 3.22. *Let x and y be nonzero elements of $\Gamma_{\mathrm{an},r}$. Then the multiplicity of a slope s of xy is the sum of its multiplicities as a slope of x and of y.*

Proof. This follows immediately from the previous lemma and the multiplicativity of the leading terms map L_s. □

Corollary 3.23. *The units of $\Gamma_{\mathrm{an,con}}$ are precisely those x with $v_n(x) = \infty$ for some n.*

Proof. A unit of $\Gamma_{\mathrm{an,con}}$ must also be a unit in $\Gamma_{\mathrm{an},r}$ for some r, and a unit of $\Gamma_{\mathrm{an},r}$ must have all slopes of multiplicity zero. (Remember, in $\Gamma_{\mathrm{an},r}$, any slopes greater than r are disregarded.) If $v_n(x) < \infty$ for all n, then x has infinitely many different slopes, so it still has slopes of nonzero multiplicity in $\Gamma_{\mathrm{an},r}$ for any r, and so can never become a unit. □

We again caution that the condition $v_n(x) = \infty$ does not imply that $x \in \Gamma_{\mathrm{con}}[\frac{1}{p}]$, if K is not finite over $k((t))$.

It will be convenient to put elements of x into a standard (multiplicative) form, so we make a statement to this effect as a lemma.

Lemma 3.24. *For any $x \in \Gamma_{\mathrm{an},r}$ nonzero, there exists a unit $u \in \Gamma_{\mathrm{an},r}$ such that ux admits a semiunit decomposition $\sum_i u_i \pi^i$ with $u_0 = 1$ and $u_i = 0$ for $i > 0$. Moreover, for such u, we have*

(a) $v_0(ux) = 0$;

(b) $w_r(ux) = 0$;

(c) $rv_n(ux) + n > 0$ *for* $n > 0$*;*

(d) the Newton polygon of ux *begins at* $(0, 0)$*.*

Proof. By Proposition 3.13, we can find a semiunit decomposition $\sum_i u'_i\pi^i$ of x; then $u'_i = 0$ for i sufficiently large. Choose the largest j such that $u'_j \neq 0$, and put $u = \pi^{-j}(u'_j)^{-1}$. Then ux admits the semiunit decomposition $\sum_i u_i\pi^i$ with $u_i = u'_{i+j}/u'_j$, so $u_0 = 1$ and $u_i = 0$ for $i > 0$.

To verify (a), note that $rv_0(ux) \geq \min_i\{rv_0(u_i\pi^i)\} \geq 0$, and the minimum is only achieved for $i = 0$: for $i < 0$, $rv_0(u_i\pi^i) > w_r(u_i\pi^i) \geq 0$ since $\sum_i u_i\pi^i$ is a semiunit decomposition of ux. Thus $rv_0(ux) = 0$, whence (a).

To verify (b), note that $w_r(ux) \geq \min_i\{w_r(u_i\pi^i)\} = 0$, whereas $w_r(ux) \leq rv_0(ux) = 0$ from (a).

To verify (c), note that for $n > 0$ and $m = v_p(\pi^i)$, $rv_n(u_i\pi^i) + n > rv_m(u_i\pi^i) + m \geq w_r(u_i\pi^i) \geq 0$, so $rv_n(ux) + n \geq \min_{i\leq 0}\{rv_n(u_i\pi^i) + n\} > 0$.

To verify (d), first note that the line through $(0, 0)$ of slope $-r$ is a lower supporting line of the set of points $(v_n(ux), n)$, since $rv_n(ux)+n \geq w_r(ux) \geq 0$ for $n \leq 0$. Thus $(0, 0)$ lies on the Newton polygon, and the slope of the segment of the Newton polygon just to the right of $(0, 0)$ is at least $-r$. We also have $rv_n(ux) + n > 0$ for $n > 0$, so the slope of the segment of the Newton polygon just to the left of $(0, 0)$, if there is one, must be less than $-r$. Thus the first segment of slope at least $-r$ does indeed begin at $(0, 0)$, as desired. □

The next lemma may be viewed as a version of the Weierstrass preparation theorem.

Lemma 3.25. *Let* x *be a nonzero element of* $\Gamma_{\mathrm{an},r}$ *whose largest slope is* s_1 *with multiplicity* $m > 0$*. Then there exists* $y \in \Gamma_{\mathrm{an},r}$*, pure of slope* s_1 *with multiplicity* m*, which divides* x*.*

Proof. If x is pure of slope s_1, there is nothing to prove. So assume that x is not pure, and let s_2 be the second largest slope of x.

By Lemma 3.24, there exists a unit $u \in \Gamma_{\mathrm{an},r}$ such that ux admits a semiunit decomposition $\sum_i u_i\pi^i$ with $u_0 = 1$ and $u_i = 0$ for $i > 0$, the slopes of x and ux occur with the same multiplicities, and the first segment of the Newton polygon of ux has left endpoint $(0, 0)$. Since that segment has slope $-s_1$ and multiplicity m, its right endpoint is $(m/s_1, -m)$. Put $M = -m/v_p(\pi)$; then $w_{s_1}(u_M\pi^M) = 0$ and $w_{s_1}(u_i\pi^i) > 0$ for $i < M$.

We first construct a sort of "Mittag-Leffler" decomposition of ux. Put $X = ux\pi^{-M}u_M^{-1}$, and set $y_0 = z_0 = 1$. Given y_l and z_l for some l, let $\sum_i w_i\pi^i$ be a semiunit decomposition of $X - y_lz_l$. Put

$$y_{l+1} = y_l + \sum_{v_0(w_i)<0} w_i\pi^i$$

$$z_{l+1} = z_l + \sum_{v_0(w_i)\geq 0} w_i\pi^i.$$

Given s with $s_2 < s < s_1$, put $c_s = w_s(X - 1)$, so that $c_s > 0$. We show that for each l, $w_s(y_l - 1) \geq c_s$, $w_s(z_l - 1) \geq c_s$, and $w_s(X - y_lz_l) \geq (l + 1)c_s$. These inequalities are clear for $l = 0$. If they hold for l, then

$$\begin{aligned} w_s(y_{l+1} - 1) &\geq \min\{w_s(y_l - 1), w_s(y_{l+1} - y_l)\} \\ &\geq \min\{c_s, (l+1)c_s\} = c_s, \end{aligned}$$

and similarly $w_s(z_{l+1} - 1) \geq c_s$. As for the third inequality, note that

$$\begin{aligned} X - y_{l+1}z_{l+1} &= X - y_l z_l + y_l(z_l - z_{l+1}) + z_{l+1}(y_l - y_{l+1}) \\ &= (y_l - 1)(z_l - z_{l+1}) + (z_{l+1} - 1)(y_l - y_{l+1}) \end{aligned}$$

since $X - y_l z_l = (y_{l+1} - y_l) + (z_{l+1} - z_l)$. Since $w_s(y_l - y_{l+1}) \geq (l+1)c_s$ and $w_s(z_l - z_{l+1}) \geq (l+1)c_s$, we conclude that

$$\begin{aligned} w_s(X - y_{l+1}z_{l+1}) &\geq \min\{w_s((y_l - 1)(z_l - z_{l+1})), w_s((z_{l+1} - 1)(y_l - y_{l+1}))\} \\ &\geq \min\{c_s + (l+1)c_s, c_s + (l+1)c_s\} \\ &= (l+2)c_s, \end{aligned}$$

as desired. This completes the induction.

We do not yet know that either $\{y_l\}$ or $\{z_l\}$ converges in $\Gamma_{\mathrm{an},r}$; to get to that point, we need to play the two sequences off of each other. Suppose s_3 satisfies $s_2 < s_3 < s_1$. Note that to get from y_l to y_{l+1}, we add terms of the form $w_i \pi^i$, with w_i a semiunit, for which $v_0(w_i) < 0$ but $sv_0(w_i) + v_p(\pi^i) \geq (l+1)c_s$ for $s_2 < s < s_1$. This implies that

$$sv_0(w_i) + v_p(\pi^i) \geq (l+1)c_{s_3}$$

for all $s \leq s_3$. In particular, $w_s(y_{l+1} - y_l) \to \infty$ as $l \to \infty$, so $\{y_l\}$ converges to a limit in $\Gamma_{\mathrm{an},s}$ for any $s \leq s_3$. Moreover, for $s \leq s_3$, we have

$$\begin{aligned} w_s(y_{l+1} - 1) &\geq \min\{w_s(y_l - 1), w_s(y_{l+1} - y_l)\} \\ &\geq \min\{w_s(y_l - 1), (l+1)c_{s_3}\} \end{aligned}$$

so by induction on l, $w_s(y_l - 1) \geq c_{s_3}$. Hence y and each of the y_l are units in Γ_{an,s_3}, for any $s_3 < s_1$.

On the flip side, to get from z_l to z_{l+1}, we add terms of the form $w_i \pi^i$, with i a semiunit, for which $v_0(w_i) > 0$ but $sv_0(w_i) + v_p(\pi^i) \geq (l+1)c$ for $s_2 < s < s_1$. This implies that $sv_0(w_i) + v_p(\pi^i) \geq (l+1)c_{s_3}$ for all $s \geq s_3$. As in the previous paragraph, we deduce $w_s(z_{l+1} - z_l) \to \infty$ and $w_s(z_l - 1) > 0$ for $s_2 < s \leq r$.

Put $z = Xy^{-1}$ in Γ_{an,s_3}. Since $w_{s_3}(y_l) = 0$ for all l, we have

$$\begin{aligned} w_{s_3}(z_l - z) &= w_{s_3}(yy_l z_l - yy_l z) \\ &= w_{s_3}(y(y_l z_l - X) + (y - y_l)X) \\ &\geq \min\{w_{s_3}(y(y_l z_l - X)), w_{s_3}((y - y_l)X)\} \end{aligned}$$

and both terms in braces tend to infinity with l. Thus $z_l \to z$ under $|\cdot|_{s_3}$.

For $s_3 \leq s \leq r$, since $sv_n(z_l - z) + n \geq (l+1)c_{s_3}$ and $sv_n(z_l) + n \to \infty$ as $n \to \pm\infty$, for any given l we have $sv_n(z) + n \geq (l+1)c_{s_3}$ for all but finitely many n. Since this holds for any l, we have $sv_n(z) + n \to \infty$ as $n \to \pm\infty$. As we already have $z \in \Gamma_{\mathrm{an},s_3}$, this is enough to imply $z \in \Gamma_{\mathrm{an},r}$. Meanwhile, put

$$a_l = X(1 + (1-z) + \cdots + (1-z)^l) = y(1 - (1-z)^{l+1}),$$

so that $w_s(a_l - y) = (l+1)w_s(1-z)$ for $s_2 < s < s_1$. In particular, for each n, $v_n(a_m - y) \to \infty$ as $m \to \infty$, and so the inequalities

$$v_n(a_l - y) \geq \min\{v_n(a_l - a_{l+1}), \dots, v_n(a_{m-1} - a_m), v_n(a_m - y)\}$$

for each m yield, in the limit as $m \to \infty$, the inequality

$$v_n(a_l - y) \geq \min\{v_n(a_l - a_{l+1}), v_n(a_{l+1} - a_{l+2}), \dots\}.$$

Now $w_s(a_{l+1} - a_l) = w_s(X(1-z)^{l+1}) = w_s(X) + (l+1)w_s(1-z)$ for $s_2 < s \leq r$, so $sv_n(a_{l+1} - a_l) + n \geq w_s(X) + (l+1)w_s(1-z)$. We conclude that

$$sv_n(a_l - y) + n \geq w_s(X) + (l+1)w_s(1-z),$$

so that $sv_n(y) + n \geq w_s(X) + (l+1)w_s(1-z)$ for all but finitely many n. Therefore $sv_n(y) + n \to \infty$ as $n \to \pm\infty$ for $s_2 < s \leq r$. Again, since we already have $y \in \Gamma_{\mathrm{an},s_3}$, we deduce that $y \in \Gamma_{\mathrm{an},r}$.

Since y is a unit in $\Gamma_{\mathrm{an},s}$ for any $s < s_1$, it has no slopes less than s_1. Since $w_s(1-z) > 0$ for $s_2 < s \leq r$, z has no slopes greater than s_2. Since the slopes of y and z together must comprise the slopes of x, y must have s_1 as a slope with multiplicity m and no other slopes, as desired. □

A *slope factorization* of a nonzero element x of $\Gamma_{\mathrm{an},r}$ is a Fréchet-convergent product $x = \prod_{j=1}^{N} x_j$ for N a positive integer or ∞, where each x_j is pure and the slopes s_j of x_j satisfy $s_1 > s_2 > \cdots$.

Lemma 3.26. *Every nonzero element of* $\Gamma_{\mathrm{an},r}$ *has a slope factorization.*

Proof. Let x be a nonzero element of $\Gamma_{\mathrm{an},r}$ with slopes $s_1, s_2, \dots$. By Lemma 3.25, we can find y_1 pure of slope s_1 dividing x such that x/y_1 has largest slope s_2. Likewise, we can find y_2 pure of slope s_2 such that y_2 divides x/y_1, y_3 pure of slope s_3 such that y_3 divides $x/(y_1y_2)$, and so on.

If there are $N < \infty$ slopes, then x and $y_1 \cdots y_N$ have the same slopes, so $x/(y_1 \cdots y_N)$ must be a unit u, and $x = (uy_1)y_2 \cdots y_N$ is a slope factorization. Suppose instead there are infinitely many slopes; then $s_i \to 0$ as $i \to \infty$. By Lemma 3.24, for each i we can find a unit a_i such that a_iy_i admits a semiunit decomposition $\sum_j u_{ij}\pi^j$ with $u_{i0} = 1$ and $u_{ij} = 0$ for

$j > 0$. For $j < 0$, $sv_n(u_{ij}\pi^j) + n$ is minimized for $n = v_p(\pi^j) < 0$ because u_{ij} is a semiunit; for i sufficiently large, we have $s \geq s_i$, so

$$\begin{aligned} sv_{v_p(\pi^j)}(u_{ij}\pi^j) + v_p(\pi^j) &= \frac{s}{s_i}(s_i v_{v_p(\pi^j)}(u_{ij}\pi^j) + v_p(\pi^j)) + \left(\frac{s}{s_i} - 1\right)(-v_p(\pi^j)) \\ &\geq \left(\frac{s}{s_i} - 1\right)(-j)v_p(\pi), \end{aligned}$$

which tends to infinity as $i \to \infty$. Hence $w_s(a_i y_i - 1) \to \infty$ as $i \to \infty$; if we put $z_j = \prod_{i=1}^{j} a_i y_i$; then $\{z_j\}$ converges to a limit z, and $\{x/z_j\}$ converges to a limit u, such that $uz = x$. The slopes of z coincide with the slopes of x, so u must be a unit, and $(ua_1y_1)\prod_{i>1}(a_i y_i)$ is a slope factorization of x. ☐

Lemma 3.27. *Let x be an element of $\Gamma_{\mathrm{an},r}$ which is pure of slope s and multiplicity m. Then for every $y \in \Gamma_{\mathrm{an},r}$, there exists $z \in \Gamma_{\mathrm{an},r}$ such that:*

(a) $y - z$ is divisible by x;

(b) $w_s(z) \geq w_s(y)$;

(c) $v_n(z) = \infty$ for $n < 0$.

Proof. Put $M = m/v_p(\pi)$. By Lemma 3.24, there exists a unit $u \in \Gamma_{\mathrm{an},r}$ such that xu admits a semiunit decomposition $\sum_{i=-M}^{0} x_i\pi^i$ with $sv_0(x_{-M}) = m$. Note that

$$w_r(x_{-M}\pi^{-M}) = rv_{-m}(x_{-M}\pi^{-M}) - m = m\left(\frac{r}{s} - 1\right).$$

Let $\sum_i y_i\pi^i$ be a semiunit decomposition of y.

We define the sequence $\{c_l\}_{l=0}^{\infty}$ of elements of $\Gamma_{\mathrm{an},r}$ such that $v_n(c_l) = \infty$ for $n < 0$, $w_r(c_l) \geq -l(v_p(\pi) + m(r/s - 1))$, $w_s(c_l) \geq -lv_p(\pi)$, and

$$c_l \equiv \pi^{-l} \pmod{x}.$$

Put $c_0 = 1$ to start. Given c_l, let $\sum_i u_i\pi^i$ be a semiunit decomposition of c_l; since $v_n(c_l) = \infty$ for $n < 0$, we have $u_i = 0$ for $i < 0$. Now set

$$c_{l+1} = \pi^{-1}(c_l - uxx_{-M}^{-1}\pi^M u_0).$$

The congruence $c_{l+1} \equiv \pi^{-1}c_l \equiv \pi^{-l-1} \pmod{x}$ is clear from the definition. Since $uxx_{-M}^{-1}\pi^M \equiv 1 \pmod{\pi}$, the term in parentheses has positive valuation, so $v_n(c_{l+1}) = \infty$ for $n < 0$. Since $w_s(ux) = w_s(x_{-M}\pi^{-M}) = 0$ and $w_s(u_0) \geq w_s(c_l)$, we have $w_s(c_{l+1}) \geq w_s(\pi^{-1}c_l) \geq -(l+1)v_p(\pi)$. Finally, $w_r(u_0) \geq w_r(c_l)$, $w_r(ux) = 0$ and $w_r(x_{-M}\pi^{-M}) = m(r/s - 1)$, so

$$\begin{aligned} w_r(c_{l+1}) &\geq w_r(\pi^{-1}) + \min\{w_r(c_l), w_r(uxx_{-M}^{-1}\pi^M u_0)\} \\ &\geq -v_p(\pi) + w_r(c_l) - m(r/s - 1) \\ &\geq -(l+1)(m(r/s - 1) + v_p(\pi)). \end{aligned}$$

We wish to show that $\sum_{i=-\infty}^{-1} y_i c_{-i}$ converges, so that its limit is congruent to $\sum_{i=-\infty}^{-1} y_i \pi^i$ modulo x. To this end, choose $t > 0$ large enough that

$$trv_p(\pi) > m(r/s - 1) + v_p(\pi).$$

Then $(1/t)v_n(y) + n \to \infty$ as $n \to -\infty$, so in particular there exists $c > 0$ such that $(1/t)v_n(y) \geq -n - c$ for $n < 0$. For $n = v_p(\pi^i)$ where $i < 0$ and $y_i \neq 0$, we have $v_n(y) = v_0(y_i)$, so we have $v_0(y_i) \geq -tiv_p(\pi) - tc$. Then

$$\begin{aligned} w_r(y_i c_{-i}) &= w_r(y_i) + w_r(c_{-i}) \\ &= rv_0(y_i) + w_r(c_{-i}) \\ &\geq -triv_p(\pi) - trc + i(m(r/s-1) + v_p(\pi)) \end{aligned}$$

which tends to infinity as $i \to -\infty$. Thus $\sum_{i=-\infty}^{-1} y_i c_{-i}$ converges under $|\cdot|_r$; since $v_n(y_i c_{-i}) = \infty$ for $n < 0$, the sum also converges under $|\cdot|_s$ for $0 < s < r$. Thus it has a limit $z' \in \Gamma_{\mathrm{an},r}$; put $z = z' + \sum_{i=0}^{\infty} y_i \pi^i$. Then $y - z = \sum_{i=-\infty}^{-1} y_i(\pi^i - c_{-i})$; since each term in the sum is divisible by x, so is the sum. This verifies (a). To verify (b), note that $w_s(y_i c_{-i}) \geq w_s(y_i \pi^i)$ for $i < 0$, so $w_s(z') \geq w_s(y)$, and clearly $w_s(z - z') \geq w_s(y)$, so $w_s(z) \geq w_s(y)$. To verify (c), simply note that each term in the sum defining z satisfies the same condition. □

3.6 The Bézout property for analytic rings

Again, we assume that the valued field K satisfies the conditions of Proposition 3.11, so that $\Gamma_{\mathrm{an,con}} = \Gamma^K_{\mathrm{an,con}}$ is defined. With the factorization results of the previous section in hand, we now focus on establishing the Bézout property for $\Gamma_{\mathrm{an,con}}$ (Theorem 3.20). We proceed by establishing principality of successively more general classes of finitely generated ideals, culminating in the desired result.

Lemma 3.28. *Let x and y be elements of $\Gamma_{\mathrm{an},r}$, each with finitely many slopes, and having no slopes in common. Then the ideal (x, y) is the unit ideal.*

Proof. We induct on the sum of the multiplicities of the slopes of x and y; the case where either x or y has total multiplicity zero is vacuous, as then x or y is a unit and so (x, y) is the unit ideal. So we assume that both x and y have positive total multiplicity.

If x is not pure, then by Lemma 3.26 it factors as $x_1 x_2$, where x_1 is pure and x_2 is not a unit. By the induction hypothesis, the ideals (x_1, y) and (x_2, y) are the unit ideal; in other words, x_1 and x_2 have multiplicative inverses modulo y. In that case, so does $x = x_1 x_2$, so (x, y) is the unit ideal. The same argument applies in case y is not pure.

It thus remains to treat the case where x and y are both pure. Let s and t be the slopes of x and y, and let m and n be the corresponding multiplicities. Put $M = m/v_p(\pi)$ and $N = n/v_p(\pi)$. Without loss of generality, we may assume $s < t$. By Lemma 3.24, we can find units u and v such that $ux = \sum_{i=-M}^{0} x_i \pi^i$ and $vy = \sum_{i=-N}^{0} y_i \pi^i$.

Put

$$X = ux\pi^M x_{-M}^{-1}, \qquad Y = vy\pi^N y_{-N}^{-1}, \qquad z = X - Y.$$

We can read off information about the Newton polygon of z by comparing $w_r(X)$ with $w_r(Y)$; see Figure 2 for an illustration. (In both diagrams, the dashed lines have slope $-r$.) If $w_r(X) < w_r(Y)$ (left side of Figure 2), then the highest vertex of the lower convex hull of the set of points $(v_l(z), l)$ occurs at $(v_m(X), m)$ and the lowest vertex has positive y-coordinate. Moreover, the slope of the first segment of the lower convex hull is at least $-s$. Thus the sum of all multiplicities of z is strictly less than m, and y and z have no common slopes, so the induction hypothesis implies that $(x, y) = (y, z)$ is the unit ideal.

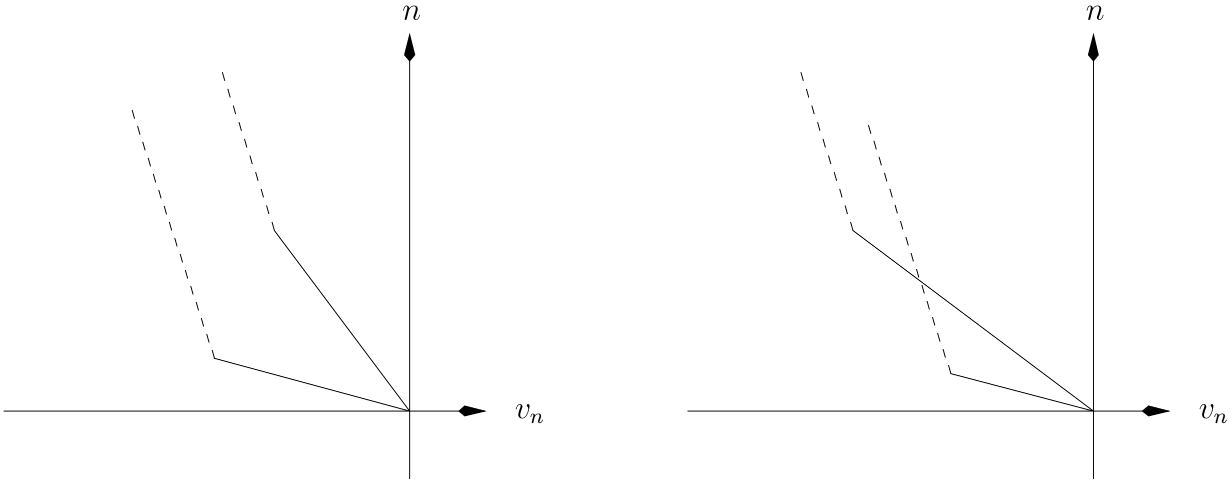

Figure 2: The Newton polygons of $X = ux\pi^M x_{-M}^{-1}$ and $Y = vy\pi^N y_{-N}^{-1}$

If $w_r(X) \geq w_r(Y)$ (right side of Figure 2), then the highest vertex of the lower convex hull of the set of points $(v_l(z), l)$ occurs at $(v_n(Y), n)$ and the lowest vertex has positive y-coordinate. Moreover, $(v_m(X), m)$ is also a vertex of the lower convex hull, and the line joining it to $(v_n(Y), n)$ is a support line of the lower convex hull. Thus the segment joining the two points is a segment of the lower convex hull, of slope less than $-t$; the remainder of the lower convex hull consists of segments of slope at least $-s$, of total multiplicity less than m. By Lemma 3.26, z factors as $z_1 z_2$, where z_1 is pure of some slope greater than t, and z_2 has all slopes less than or equal to s and total multiplicity less than m. By the induction hypothesis, (x, z_1) and (y, z_2) both equal the unit ideal. But $(y, z_1) = (x, z_1)$ since z_1 divides $z = ux\pi^M x_{-M}^{-1} - vy\pi^N y_{-N}^{-1}$, so $(y, z_1 z_2) = (y, z) = (x, y)$ is also equal to the unit ideal.

We conclude that the induction goes through for all x and y. This completes the proof. $\square$

Lemma 3.29. *Let x and y be elements of $\Gamma_{\mathrm{an},r}$ with x, y pure of the same slope s. Then (x, y) is either the unit ideal or is generated by a pure element of slope s.*

Proof. Let m and n be the multiplicities of s as a slope of x and y; we induct on $m + n$. Put $M = m/v_p(\pi)$ and $N = n/v_p(\pi)$. By Lemma 3.24, we may choose units u, v so that ux and vy admit semiunit decompositions $ux = \sum_{i=-M}^0 x_i \pi^i$ and $vy = \sum_{i=-N}^0 y_i \pi^i$. Put

$$X = ux\pi^M x_{-M}^{-1}, \qquad Y = vy\pi^N y_{-N}^{-1}, \qquad z = X - Y.$$

By symmetry, we may assume $m \leq n$ without loss of generality.

First suppose $m < n$. Then the highest vertex of the lower convex hull of the set of points $(v_l(z), l)$ occurs at $(v_n(Y), n)$, and the lowest vertex has positive y-coordinate. Thus the total multiplicity of z is strictly less than n. By Lemma 3.26, we may write $z = z_1 z_2$, where z_1 has no slopes equal to s, and z_2 is either a unit or is pure of slope s. By Lemma 3.28, (x, z_1) is the unit ideal, so $(x, y) = (x, z) = (x, z_2)$, which is principal by the induction hypothesis.

Next suppose $m = n$. We define a sequence $\{z_l\}$ as follows, starting with $z_0 = z = X - Y$. Given z_l, let $\sum_i u_i \pi^i$ be a semiunit decomposition of z_l, and set

$$z_{l+1} = z_l - \sum u_i \pi^i Y$$

taking the sum over all i for which $w_r(u_i \pi^i) > 0$. Clearly $z_l - z_0$ is divisible by y for each l.

Suppose $z_l = z_{l+1}$ for some l. If $\sum_i u_i \pi^i$ is the chosen semiunit decomposition of z_l, then $w_r(u_i \pi^i) \leq 0$ for each i. On the other hand, we have $w_r(z_l - z_0) > w_r(Y)$ by induction on l.

We claim that the multiplicity of s as a slope of z_l is less than n. Suppose that this is not the case; then there exist indices $i < j$ with $w_r(u_i \pi^i) - w_r(u_j \pi^j) = -w_r(Y)$. As shown above, we have $w_r(u_i \pi^i) \leq 0$; since $w_r(z_l - z_0) > w_r(Y)$ and $w_r(z_0) \geq w_r(Y)$, we also have $w_r(u_j \pi^j) \geq w_r(Y)$. Since these inequalities imply $w_r(u_i \pi^i) - w_r(u_j \pi^j) \leq -w_r(Y)$, they must all be equalities; that is, $w_r(u_i \pi^i) = 0$, $w_r(u_j \pi^j) = w_r(Y)$, and $w_r(z_0) = w_r(Y)$. This means that the highest vertex of the lower convex hull of the set of points $(v_c(z_l), c)$ occurs at $(v_n(Y), n)$, and there is another vertex with y-coordinate at most 0. This is a contradiction.

Still under the hypothesis $z_l = z_{l+1}$, we conclude that the multiplicity of s as a slope of z_l is less than n. By Lemma 3.26, we can write $z_l = a_1 a_2$, where a_1 has no slopes equal to s, and a_2 is pure of slope s and multiplicity less than n. By Lemma 3.28, (y, a_1) is the unit ideal; by the induction hypothesis, (y, a_2) is principal. Thus $(x, y) = (y, z_0) = (y, z_l) = (y, a_1 a_2) = (y, a_2)$ is principal, as desired.

Finally, suppose that $z_l \neq z_{l+1}$ for all l. We claim that $z_l \equiv 0 \pmod{\pi^{l+1}}$ for each l. This holds for $l = 0$ by the choice of z_0. Given the congruence for z_l, let $\sum_i u_i \pi^i$ be the chosen semiunit decomposition of z_l. By hypothesis, $u_i = 0$ for $i \leq l$. If $u_{l+1} \neq 0$, then we must have $w_r(u_{l+1} \pi^{l+1}) > 0$, or else $w_r(u_i \pi^i) \leq 0$ for all i (recall that in a semiunit decomposition, whenever $i < j$ and $u_i, u_j \neq 0$ we have $w_r(u_i \pi^i) \geq w_r(u_j \pi^j)$). Thus $z_{l+1} \equiv z_l - u_{l+1} \pi^{l+1} \equiv 0 \pmod{\pi^{l+2}}$. Thus the congruence holds by induction.

Since the z_l converge to zero π-adically and $w_r(z_l)$ is bounded below, the z_l also converge to zero in the Fréchet topology. This means the series $\sum_{l=0}^{\infty} (z_l - z_{l+1})$ converges to z_0, the series $\sum_{l=0}^{\infty} (z_l - z_{l+1})/y$ converges to some limit a, and we have $ya = z_0$. Therefore z_0 is divisible by y, as is x, and the ideal (x, y) is generated by y.

This completes the induction in all cases, whence the desired result. □

Corollary 3.30. *For $x, y \in \Gamma_{\mathrm{an},r}$ with x pure of slope s, the ideal (x, y) is principal.*

Proof. By Lemma 3.27, there exists $z \in \Gamma_{\mathrm{an},r}$ such that $y - z$ is divisible by x and $v_n(z) = \infty$ for $n < 0$. Thus z has only finitely many slopes. By Lemma 3.26, we can factor z as $z_1 z_2$, where z_1 is pure of slope s and z_2 has no slopes equal to s. Then (x, z_2) is the unit ideal, so $(x, y) = (x, z) = (x, z_1)$, which is principal by Lemma 3.29. □

Lemma 3.31 (Principal parts theorem). *Let s_n be a decreasing sequence of positive rationals with limit* 0*, and suppose $x_n \in \Gamma_{\mathrm{an},r}$ is pure of slope s_n for all n. Then for any sequence y_n of elements of $\Gamma_{\mathrm{an},r}$, there exists $y \in \Gamma_{\mathrm{an},r}$ such that $y \equiv y_n \pmod{x_n}$ for all n.*

Proof. As in the proof of Lemma 3.26, we can replace each x_n with itself times a unit, in such a way that $\prod_n x_n$ converges. Put $x = \prod_n x_n$ and $u_n = x/x_n$. By Lemma 3.29, x_n is coprime to each of $x_1, \dots, x_{n-1}$. By Corollary 3.30, the ideal $(x_n, \prod_{i>n} x_i)$ is principal, but if it were not the unit ideal, any generator would both be pure of slope s_n and have all slopes less than s_n. Thus x_n is coprime to $\prod_{i>n} x_i$, hence also to u_n.

We construct a sequence $\{z_n\}_{n=1}^{\infty}$ such that $u_n z_n \equiv y_n \pmod{x_n}$ and $\sum u_n z_n$ converges for the Fréchet topology; then we may set $y = \sum u_n z_n$ and be done. For the moment, fix n and choose v_n with $u_n v_n \equiv y_n \pmod{x_n}$.

For $s > s_n$, we have $|1 - x_n|_s < 1$, so the sequence $c_m = -1 - (1 - x_n) - \cdots - (1 - x_n)^m$ is Cauchy for the norm $|\cdot|_s$, and $|1 + c_m x_n|_s = |1 - x_n|_s^{m+1} \to 0$ under $|\cdot|_s$. In particular, for any $\epsilon > 0$, there exists m such that $|1 + c_m x_n|_s < \epsilon$ for $s_{n-1} \leq s \leq r$.

Now choose $\epsilon_n > 0$ such that $\epsilon_n |u_n v_n|_s < 1/n$ for all $s \geq s_{n-1}$ (with n still fixed), choose m as above, and put $z_n = v_n(1 + c_m x_n)$. Then $u_n z_n \equiv u_n v_n \equiv y_n \pmod{x_n}$. Moreover, for any $s > 0$, we have $s \geq s_{n-1}$ for sufficiently large n since the s_n tend to zero. Thus for n sufficiently large,

$$\begin{aligned} |u_n z_n|_s &= |u_n v_n (1 + c_m x_n)|_s \\ &< \epsilon_n |u_n v_n|_s < 1/n. \end{aligned}$$

Hence $\sum_n u_n z_n$ converges with respect to $|\cdot|_s$ for $0 < s \leq r$, and its limit y has the desired property. ☐

At long last, we are ready to prove the generalization of Lazard's result, that $\Gamma_{\mathrm{an},r}$ is a Bézout ring.

Proof of Theorem 3.20. By induction on the number of generators of the ideal, it suffices to prove that if $x, y \in \Gamma_{\mathrm{an},r}$ are nonzero, then the ideal (x, y) is principal.

Pick a slope factorization $\prod_j y_j$ of y. By Corollary 3.30, we can choose a generator d_j of (x, y_j) for each j, such that d_j is either 1 or is pure of the same slope as y_j. As in the proof of Lemma 3.26, we can choose the d_j so that $\prod_j d_j$ converges. Since the d_j are pairwise coprime by Lemma 3.28, x is divisible by the product of any finite subset of the d_j, and hence by $\prod_j d_j$.

Choose a_j and b_j such that $a_j x + b_j y_j = d_j$, and apply Lemma 3.31 to find z such that $z \equiv a_j \prod_{k \neq j} d_j \pmod{y_j}$ for each j. Then $zx - \prod_j d_j$ is divisible by each y_j, so it is divisible by y, and so $\prod_j d_j$ generates the ideal (x, y). Thus (x, y) is principal and the proof is complete. ☐

Corollary 3.32. *For K a finite extension of $k((t))$, the ring $\Gamma_r^K[\frac{1}{p}]$ is a Bézout ring.*

Proof. For $x, y \in \Gamma_r^K[\frac{1}{p}]$, Theorem 3.20 implies that the ideal (x, y) becomes principal in $\Gamma_{\mathrm{an},r}^K$. Let d be a generator; then d must have finite total multiplicity, and so belongs to $\Gamma_r^K[\frac{1}{p}]$.

Put $x' = x/d$ and $y' = y/d$, so that (x', y') becomes the unit ideal in $\Gamma_{\mathrm{an},r}^K$. By Lemma 3.26, x' factors in $\Gamma_{\mathrm{an},r}^K$ as $a_1 \dots a_l$, where each a_i is pure. Since each of those factors has finite total multiplicity, each lies in $\Gamma_r^K[\frac{1}{p}]$.

Since (x', y') is the unit ideal in $\Gamma_{\mathrm{an},r}^K$, so is (a_i, y') for each i. That is, there exist b_i and c_i in $\Gamma_{\mathrm{an},r}^K$ such that $a_i b_i + c_i y' = 1$. Since a_i is pure, Lemma 3.27 implies that $c_i \equiv d_i \pmod{a_i}$ for some d_i with finite total multiplicity, which thus belongs to $\Gamma_r^K[\frac{1}{p}]$. Now $d_i y' \equiv 1 \pmod{a_i}$, and $e_i = (d_i y' - 1)/a_i$ has finite total multiplicity, so itself lies in $\Gamma_r^K[\frac{1}{p}]$. We now have the relation $a_i e_i + d_i y' = 1$ within $\Gamma_r^K[\frac{1}{p}]$, so (a_i, y') is the unit ideal in $\Gamma_r^K[\frac{1}{p}]$. Since this is true for each i, (x', y') is also the unit ideal and so $(x, y) = (d)$.

We conclude that any ideal generated by two elements is principal. By induction, this implies that $\Gamma_r^K[\frac{1}{p}]$ has the Bézout property. □

One presumably has the same result if K is perfect, but it does not follow formally from Theorem 3.20, since Γ_r^K is not Fréchet complete in $\Gamma_{\mathrm{an},r}^K$. That is, an element of $\Gamma_{\mathrm{an},r}^K$ of finite total multiplicity need not lie in Γ_r^K. So one must repeat the arguments used to prove Theorem 3.20 working withing $\Gamma_r^K[\frac{1}{p}]$; as we have no use for the result, we leave this to the reader.

4 The special Newton polygon

In this chapter, we construct a Newton polygon for σ-modules over $\Gamma_{\mathrm{an,con}}$, the "special Newton polygon". More precisely, we give a slope filtration over $\Gamma_{\mathrm{an,con}}^{\mathrm{alg}}$ that, in case the σ-module is quasi-unipotent, is precisely the filtration that makes it quasi-unipotent. The special Newton polygon is a numerical invariant of this filtration.

Throughout this chapter, we assume K is a valued field satisfying the condition of Proposition 3.11. The choice of K will only be relevant once or twice, as most of the time we will be working with $\Gamma_{\mathrm{an,con}}^{\mathrm{alg}} = \Gamma_{\mathrm{an,con}}^{k((t))^{\mathrm{alg}}}$. When this is the case, we will also assume k is algebraically closed and that $\pi^\sigma = \pi$.

We will use without further comment the facts that every element of $\Gamma_{\mathrm{an,con}}^{\mathrm{alg}}$ has a strong semiunit decomposition (Proposition 3.14) and that $\Gamma_{\mathrm{an,con}}$ and $\Gamma_{\mathrm{an,con}}^{\mathrm{alg}}$ are Bézout rings (Theorem 3.20). In particular, any σ-module over $\Gamma_{\mathrm{an,con}}$ or $\Gamma_{\mathrm{an,con}}^{\mathrm{alg}}$ is free by Proposition 2.5, so admits a basis.

4.1 Properties of eigenvectors

Recall that we call a nonzero element $\mathbf{v}$ of a σ-module M an *eigenvector* if there exists $\lambda \in \mathcal{O}[\frac{1}{p}]$ such that $F\mathbf{v} = \lambda \mathbf{v}$. Also recall that if $\mathbf{v}$ an eigenvector, the *slope* of $\mathbf{v}$ is defined to be $v_p(\lambda)$. (Beware: this differs from the notion of slope used in Section 3.5.) Our method of constructing the special Newton polygon of a σ-module over $\Gamma_{\mathrm{an,con}}^{\mathrm{alg}}$ is to exhibit a basis of eigenvectors after enlarging $\mathcal{O}$ suitably. Before proceeding, it behooves us to catalog some basic properties of eigenvectors of σ-modules over $\Gamma_{\mathrm{an,con}}^{\mathrm{alg}}$. Some of these assertions will also hold more generally over $\Gamma_{\mathrm{an,con}}$ (i.e., for arbitrary K), so we distinguish between $\Gamma_{\mathrm{an,con}}$ and $\Gamma_{\mathrm{an,con}}^{\mathrm{alg}}$ in the statements below.

For M a σ-module over $\Gamma_{\mathrm{an,con}}$, we say $\mathbf{v} \in M$ is *primitive* if $\mathbf{v}$ extends to a basis of M. By Lemma 2.3, if $\mathbf{e}_1, \dots, \mathbf{e}_n$ is a basis of M and $\mathbf{v} = \sum c_i \mathbf{e}_i$, then $\mathbf{v}$ is primitive if and only if the c_i generate the unit ideal.

Lemma 4.1. *Let M be a σ-module over $\Gamma_{\mathrm{an,con}}^{\mathrm{alg}}$. Then every eigenvector of M is a multiple of a primitive eigenvector.*

Proof. Suppose $F\mathbf{v} = \lambda \mathbf{v}$. Choose a basis $\mathbf{e}_1, \dots, \mathbf{e}_n$, put $\mathbf{v} = \sum_i c_i \mathbf{e}_i$, and let I be the ideal generated by the c_i. Then I is invariant under σ and σ^{-1}. By Theorem 3.20, I is principal; if r is a generator of I, then r^σ is also a generator. Put $r^\sigma = cr$, with c a unit, and write $c = \mu d$, with $\mu \in \mathcal{O}[\frac{1}{p}]$, $v_0(d) < \infty$ and $v_n(d) = \infty$ for $n < 0$. By Proposition 3.18, there exists a unit $s \in \Gamma_{\mathrm{an,con}}^{\mathrm{alg}}$ such that $s^\sigma = ds$; then $(r/s)^\sigma = \mu(r/s)$. Therefore $\sum_i s(c_i/r)\mathbf{e}_i$ is a primitive eigenvector of M of which $\mathbf{v}$ is a multiple, as desired. ☐

A sort of converse to the previous statement is the following.

Proposition 4.2. *For M a σ-module over $\Gamma_{\mathrm{an,con}}^{\mathrm{alg}}$, if M contains an eigenvector of eigenvalue $\lambda \in \mathcal{O}[\frac{1}{p}]$, then it contains an eigenvector of eigenvalue $\lambda\mu$ for any $\mu \in \mathcal{O}$.*

Proof. Let $\mathbf{v} \in M$ be a nonzero eigenvector with $F\mathbf{v} = \lambda \mathbf{v}$. If μ is a unit, there exists a unit $c \in \mathcal{O}$ such that $c^\sigma = \mu c$. If μ is not a unit, then by Proposition 3.19(b) there exists a nonzero $c \in \Gamma_{\mathrm{an,con}}^{\mathrm{alg}}$ such that $c^\sigma = \mu c$. In either case, we have $F(c\mathbf{v}) = c^\sigma \lambda \mathbf{v} = \lambda\mu(c\mathbf{v})$. ☐

Proposition 4.3. *Let $0 \to M_1 \to M \to M_2 \to 0$ be an exact sequence of σ-modules over $\Gamma_{\mathrm{an,con}}^{\mathrm{alg}}$. Assume M_1 and M_2 have bases $\mathbf{v}_1, \dots, \mathbf{v}_m$ and $\mathbf{w}_1, \dots, \mathbf{w}_n$ of eigenvectors such that the slope of $\mathbf{v}_i$ is less than or equal to the slope of $\mathbf{w}_j$ for $1 \leq i \leq m$ and $1 \leq j \leq n$. Then the exact sequence splits over $\Gamma_{\mathrm{an,con}}^{\mathrm{alg}}$.*

Proof. Choose a basis $\mathbf{v}_1, \dots, \mathbf{v}_m, \mathbf{x}_1, \dots, \mathbf{x}_n$ of M such that $\mathbf{x}_j$ projects onto $\mathbf{w}_j$ in M_2 for $j = 1, \dots, n$. Suppose $F\mathbf{v}_i = \lambda_i \mathbf{v}_i$ for some $\lambda_i \in \mathcal{O}[\frac{1}{p}]$. Then $F\mathbf{x}_j = \mu_j \mathbf{x}_j + \sum_{i=1}^m A_{ij}\mathbf{v}_i$ for some $\mu_j \in \mathcal{O}[\frac{1}{p}]$ and $A_{ij} \in \Gamma_{\mathrm{an,con}}^{\mathrm{alg}}$. If $\mathbf{y}_j = \mathbf{x}_j + \sum_{i=1}^m c_{ij}\mathbf{v}_i$, then

$$F\mathbf{y}_j = \mu_j \mathbf{y}_j + \sum_{i=1}^m (\lambda_i c_{ij}^\sigma - \mu_j c_{ij} + A_{ij})\mathbf{v}_i.$$

By Proposition 3.19(a) and (b), we can choose $c_{ij} \in \Gamma_{\mathrm{an,con}}^{\mathrm{alg}}$ for each i, j so that $\lambda_i c_{ij}^\sigma - \mu_j c_{ij} + A_{ij} = 0$. For this choice, $\mathbf{v}_1, \dots, \mathbf{v}_m, \mathbf{y}_1, \dots, \mathbf{y}_n$ form a basis of eigenvectors, so the exact sequence splits as desired. ☐

Proposition 4.4. *Let M be a σ-module over $\Gamma_{\mathrm{an,con}}$ with a basis $\mathbf{w}_1, \dots, \mathbf{w}_n$ such that $F\mathbf{w}_i = \mu_i \mathbf{w}_i + \sum_{j<i} A_{ij}\mathbf{w}_j$ for some $\mu_i \in \mathcal{O}[\frac{1}{p}]$ and $A_{ij} \in \Gamma_{\mathrm{an,con}}$. Then any eigenvector of M has slope at least $\min_i\{v_p(\mu_i)\}$.*

Proof. Let $\mathbf{v}$ be any eigenvector of M, with $F\mathbf{v} = \lambda\mathbf{v}$. Write $\mathbf{v} = \sum_i b_i\mathbf{w}_i$ for some $b_i \in \Gamma_{\mathrm{an,con}}$. Suppose that $v_p(\lambda) < v_p(\mu_i)$ for all i. Then

$$\sum_i \lambda b_i \mathbf{w}_i = F\mathbf{v} = \sum_i b_i^\sigma \mu_i \mathbf{w}_i + \sum_i b_i^\sigma \sum_{j<i} A_{ij}\mathbf{w}_j.$$

Comparing the coefficients of $\mathbf{w}_n$ yields $\lambda b_n = \mu_n b_n^\sigma$, which implies $b_n = 0$ by Proposition 3.19(c). Then comparing the coefficients of $\mathbf{w}_{n-1}$ yields $\lambda b_{n-1} = \mu_{n-1} b_{n-1}^\sigma$, so $b_{n-1} = 0$. Continuing in this fashion, we deduce $b_1 = \cdots = b_n = 0$, contradiction. Thus $v_p(\lambda) \geq v_p(\mu_i)$ for some i, as desired. □

Recall that a sequence $(a_1, \dots, a_n)$ of real numbers is said to *majorize* another sequence $(b_1, \dots, b_n)$ if $a_1 + \cdots + a_n = b_1 + \cdots + b_n$ and for $i = 1, \dots, n-1$, the sum of the i smallest of $a_1, \dots, a_n$ is less than or equal to the sum of the i smallest of $b_1, \dots, b_n$. Note that two sequences majorize each other if and only if they are equal up to permutation.

Proposition 4.5. *Let M be a σ-module over $\Gamma_{\mathrm{an,con}}$ with a basis $\mathbf{v}_1, \dots, \mathbf{v}_n$ of eigenvectors, with $F\mathbf{v}_i = \lambda_i\mathbf{v}_i$ for $\lambda_i \in \mathcal{O}[\frac{1}{p}]$. Let $\mathbf{w}_1, \dots, \mathbf{w}_n$ be a basis of M such that $F\mathbf{w}_i = \mu_i\mathbf{w}_i + \sum_{j<i} A_{ij}\mathbf{w}_j$ for some $\mu_i \in \mathcal{O}[\frac{1}{p}]$ and $A_{ij} \in \Gamma_{\mathrm{an,con}}$. Then the sequence $v_p(\mu_1), \dots, v_p(\mu_n)$ majorizes the sequence $v_p(\lambda_1), \dots, v_p(\lambda_n)$.*

Proof. Assume without loss of generality that $v_p(\lambda_1) \geq \cdots \geq v_p(\lambda_n)$. Note that $v_p(\mu_1) + \cdots + v_p(\mu_n) = v_p(\lambda_1) + \cdots + v_p(\lambda_n)$ since both are equal to the slopes of primitive eigenvectors of $\wedge^n M$. Note also that $\wedge^i M$ satisfies the conditions of Proposition 4.4 for all i, using the exterior products of the $\mathbf{w}_j$ as the basis and the corresponding products of the μ_j as the diagonal matrix entries. (More precisely, view the exterior products as being partially ordered by sum of indices; any total ordering of the products refining this partial order satisfies the conditions of the proposition.) Since $\mathbf{v}_{n-i+1} \wedge \cdots \wedge \mathbf{v}_n$ is an eigenvector of $\wedge^i M$ of slope $v_p(\lambda_{n-i+1}) + \cdots + v_p(\lambda_n)$, by Proposition 4.4 this slope is greater than or equal to the smallest valuation of an i-term product of the μ_j, i.e., the sum of the i smallest of $v_p(\mu_1), \dots, v_p(\mu_n)$. This is precisely the desired majorization. □

Corollary 4.6. *Let M be a σ-module over $\Gamma_{\mathrm{an,con}}$. If $\mathbf{v}_1, \dots, \mathbf{v}_n$ and $\mathbf{w}_1, \dots, \mathbf{w}_n$ are bases of M such that $F\mathbf{v}_i = \lambda_i\mathbf{v}_i$ and $F\mathbf{w}_i = \mu_i\mathbf{w}_i$ for some $\lambda_i, \mu_i \in \mathcal{O}[\frac{1}{p}]$, then the sequences $v_p(\lambda_1), \dots, v_p(\lambda_n)$ and $v_p(\mu_1), \dots, v_p(\mu_n)$ are permutations of each other.*

Finally, we observe that the existence of an eigenvector of a specified slope does not depend on what ring of scalars $\mathcal{O}$ is used, so long as the value group of $\mathcal{O}$ contains the desired slope.

Proposition 4.7. *Let M be a σ-module over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$. Suppose $\lambda \in \mathcal{O}[\frac{1}{p}]$ occurs as the eigenvalue of an eigenvector of $M \otimes_{\mathcal{O}} \mathcal{O}'$ for some finite extension $\mathcal{O}'$ of $\mathcal{O}$. Then λ occurs as the eigenvalue of an eigenvector of M.*

Proof. Since k here is algebraically closed, we can choose a basis $\mu_1, \dots, \mu_m$ of $\mathcal{O}'$ over $\mathcal{O}$ consisting of elements fixed by σ. (Namely, let π' be a uniformizer of $\mathcal{O}'$ fixed by σ, and take $\mu_i = (\pi')^{i-1}$.) Given an eigenvector $\mathbf{v}$ over $\mathcal{O}'[\frac{1}{p}]$ with $F\mathbf{v} = \lambda\mathbf{v}$, we can write $\mathbf{v} = \mu_1\mathbf{w}_1 + \cdots + \mu_m\mathbf{w}_m$ for a unique choice of $\mathbf{w}_1, \dots, \mathbf{w}_m \in M$. Now

$$0 = F\mathbf{v} - \lambda\mathbf{v} = \mu_1(F\mathbf{w}_1 - \lambda\mathbf{w}_1) + \cdots + \mu_m(F\mathbf{w}_m - \mathbf{w}_m).$$

Since the representation $0 = \mu_1(0) + \cdots + \mu_m(0)$ is unique, we must have $F\mathbf{w}_i = \lambda\mathbf{w}_i$ for $i = 1, \dots, m$. Since $\mathbf{v}$ is nonzero, at least one of the $\mathbf{w}_i$ must be nonzero, and it provides the desired eigenvector within M. ☐

4.2 Existence of eigenvectors

In this section, we prove that every σ-module over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ has an eigenvector.

Proposition 4.8. *For every σ-module M over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$, there exist $\lambda \in \mathcal{O}_0$ and $\mathbf{v} \in M$, both nonzero, such that $F\mathbf{v} = \lambda\mathbf{v}$.*

Note that once this assertion is established for a single λ, it holds for all $\lambda \in \mathcal{O}$ of sufficiently high valuation by Proposition 3.19(b).

Proof. Let v denote the valuation on $k((t))^{\mathrm{alg}}$ normalized so that $v(t) = 1$. Let $\mathbf{e}_1, \dots, \mathbf{e}_n$ be a basis for M, and suppose $F\mathbf{e}_i = \sum_j A_{ij}\mathbf{e}_j$. Choose $r > 0$ so that the entries of A_{ij} all lie in $\Gamma^{\mathrm{alg}}_{\mathrm{an},r}$, and let c be an integer less than $\min\{w_r(A), w_r((A^{-1})^{\sigma^{-1}})\}$. For $0 < s \leq r$, we define the valuations w_s on M in terms of the basis $\mathbf{e}_1, \dots, \mathbf{e}_n$. That is, $w_s(\sum_i c_i\mathbf{e}_i) = \min_i\{w_s(c_i)\}$.

Notice that for $\lambda \in \mathcal{O}_0$ and u a strong semiunit,

$$\begin{aligned}
v_0(u) \geq \frac{-c + v_p(\lambda)}{(q-1)r} &\iff v_0(u)r \leq -v_p(\lambda) + v_0(u)qr + c \\
&\implies w_r(u\mathbf{e}_i) < w_r(\lambda^{-1}F(u\mathbf{e}_i)) \\
v_0(u) \leq \frac{qc + qv_p(\lambda)}{(q-1)r} &\iff v_0(u)r \leq v_p(\lambda) + v_0(u)r/q + c \\
&\implies w_r(u\mathbf{e}_i) < w_r(\lambda F^{-1}(u\mathbf{e}_i)).
\end{aligned}$$

Choose $\lambda \in \mathcal{O}_0$ of large enough valuation so that $-c + v_p(\lambda) < qc + qv_p(\lambda)$, and let d be a rational number such that $d(q-1)r \in (-c + v_p(\lambda), qc + qv_p(\lambda))$.

Define functions $a, b, f : M \to M$ as follows. Given $\mathbf{w} \in M$, write $\mathbf{w} = \sum_{i=1}^n z_i\mathbf{e}_i$, let $z_i = \sum_m \pi^m u_{i,m}$ be a strong semiunit decomposition for each i, let x_i be the sum of $\pi^m u_{i,m}$ over all m such that $v_0(u_{i,m}) < d$, and put $y_i = z_i - x_i$. Put $a(\mathbf{w}) = \sum_{i=1}^n x_i\mathbf{e}_i$, $b(\mathbf{w}) = \sum_{i=1}^n y_i\mathbf{e}_i$, and

$$f(\mathbf{w}) = \lambda^{-1}b(\mathbf{w}) - F^{-1}a(\mathbf{w}).$$

(Note: the definitions of a, b, f depend on the choices of semiunit decompositions above, but this does not cause any trouble.) From the inequalities tabulated above, we have

$$w_r(\lambda F^{-1}a(\mathbf{w})) \geq w_r(a(\mathbf{w})) + \epsilon, \qquad w_r(\lambda^{-1}Fb(\mathbf{w})) \geq w_r(b(\mathbf{w})) + \epsilon$$

for some $\epsilon > 0$. Therefore

$$\begin{aligned} w_r(f(\mathbf{w})) &= w_r(\lambda^{-1}b(\mathbf{w}) - F^{-1}a(\mathbf{w})) \\ &\geq w_r(\lambda^{-1}\mathbf{w}) \\ w_r(F(f(\mathbf{w})) - \lambda f(\mathbf{w}) + \mathbf{w}) &= w_r(F\lambda^{-1}b(\mathbf{w}) - a(\mathbf{w}) - b(\mathbf{w}) - \lambda F^{-1}a(\mathbf{w}) + \mathbf{w}) \\ &= w_r(\lambda^{-1}Fb(\mathbf{w}) + \lambda F^{-1}a(\mathbf{w})) \\ &\geq w_r(\mathbf{w}) + \epsilon \end{aligned}$$

for all nonzero $\mathbf{w} \in M$.

Now define sequences $\{\mathbf{v}_l\}_{l=0}^{\infty}$ and $\{\mathbf{w}_l\}_{l=0}^{\infty}$ as follows. First choose $T \in k((t))^{\mathrm{alg}}$ of valuation d, and set

$$\mathbf{v}_0 = \lambda^{-1}[T]\mathbf{e}_1 + [T^{1/q}]F^{-1}\mathbf{e}_1,$$

where the brackets again denote Teichmüller lifts. Then define $\mathbf{v}_l$ and $\mathbf{w}_l$ recursively by the formulas

$$\mathbf{w}_l = F\mathbf{v}_l - \lambda \mathbf{v}_l, \qquad \mathbf{v}_{l+1} = \mathbf{v}_l + f(\mathbf{w}_l).$$

For each l, $\mathbf{v}_l$ is defined over $\Gamma^{\mathrm{alg}}_{\mathrm{an},rq}$ and $\mathbf{w}_l$ is defined over $\Gamma^{\mathrm{alg}}_{\mathrm{an},r}$. By the final remark of the previous paragraph, we have

$$w_r(\mathbf{v}_{l+1} - \mathbf{v}_l) = w_r(f(\mathbf{w}_l)) \geq w_r(\lambda^{-1}\mathbf{w}_l)$$

and

$$\begin{aligned} w_r(\mathbf{w}_l) &= w_r(F\mathbf{v}_l - \lambda\mathbf{v}_l) \\ &= w_r(F\mathbf{v}_{l-1} + Ff(\mathbf{w}_{l-1}) - \lambda\mathbf{v}_{l-1} - \lambda f(\mathbf{w}_{l-1})) \\ &= w_r(Ff(\mathbf{w}_{l-1}) - \lambda f(\mathbf{w}_{l-1}) + \mathbf{w}_{l-1}) \\ &\geq w_r(\mathbf{w}_{l-1}) + \epsilon. \end{aligned}$$

Thus $w_r(\mathbf{w}_l)$ is a strictly increasing function of l that tends to ∞, and $w_r(\mathbf{v}_{l+1} - \mathbf{v}_l)$ also tends to ∞ with l.

We claim that in the Fréchet topology, $\mathbf{w}_l$ converges to 0 and so $\mathbf{v}_l$ converges to a limit $\mathbf{v}$, from which it follows that $F\mathbf{v} - \lambda\mathbf{v} = \lim_{l\to\infty} \mathbf{w}_l = 0$. We first show that $w_s(\lambda F^{-1}a(\mathbf{w}_l)) \to \infty$ as $l \to \infty$ for $0 < s \leq qr$.

Let $a(\mathbf{w}_l) = \sum_{i,m} \pi^m a_{l,i,m}\mathbf{e}_i$ be a strong semiunit decomposition, in which we must have $v_0(a_{l,i,m}) < d$ whenever $a_{l,i,m} \neq 0$. Then

$$\begin{aligned} w_s(\lambda F^{-1}a(\mathbf{w}_l)) &\geq w_s(\lambda(A^{-1})^{\sigma^{-1}}) + w_s(a(\mathbf{w}_l)^{\sigma^{-1}}) \\ &= w_s(\lambda(A^{-1})^{\sigma^{-1}}) + w_{s/q}(a(\mathbf{w}_l)) \\ &= w_s(\lambda(A^{-1})^{\sigma^{-1}}) + \min_{i,m}\{mv_p(\pi) + (s/q)v_0(a_{l,i,m})\} \\ &\geq w_s(\lambda(A^{-1})^{\sigma^{-1}}) + \min_{i,m}\{mv_p(\pi) + rv_0(a_{l,i,m})\} + \min_{i,m}\{(-r + s/q)v_0(a_{l,i,m})\} \\ &> w_s(\lambda(A^{-1})^{\sigma^{-1}}) + w_r(a(\mathbf{w}_l)) - (r - s/q)d. \end{aligned}$$

In particular, $w_s(\lambda F^{-1}a(\mathbf{w}_l)) \to \infty$ as $l \to \infty$.

We next show that $w_s(\lambda^{-1}Fb(\mathbf{w}_l)) \to \infty$ as $l \to \infty$ for $0 < s \leq r$. Let $b(\mathbf{w}_l) = \sum_{i,m} \pi^m b_{l,i,m}\mathbf{e}_i$ be a strong semiunit decomposition, necessarily with $v_0(b_{l,i,m}) \geq d$ whenever $b_{l,i,m} \neq 0$. Then

$$\begin{aligned} w_s(\lambda^{-1}Fb(\mathbf{w}_l)) &\geq w_s(\lambda^{-1}A) + w_s(b(\mathbf{w}_l)^\sigma) \\ &= w_s(\lambda^{-1}A) + w_{sq}(b(\mathbf{w}_l)) \\ &= w_s(\lambda^{-1}A) + \min_{i,m}\{mv_p(\pi) + sqv_0(b_{l,i,m})\}. \end{aligned}$$

Choose $e > 0$ large enough so that $s(q-1)e + w_s(\lambda^{-1}A) > 0$. If $v_0(b_{l,i,m}) < e$, then

$$\begin{aligned} mv_p(\pi) + sqv_0(b_{l,i,m}) &= mv_p(\pi) + rv_0(b_{l,i,m}) + (sq-r)v_0(b_{l,i,m}) \\ &\geq w_r(b(\mathbf{w}_l)) + h, \end{aligned}$$

where $h = (sq-r)d$ if $sq - r \geq 0$ and $h = (sq-r)e$ if $sq - r < 0$. If $v_0(b_{l,i,m}) \geq e$, then

$$\begin{aligned} mv_p(\pi) + sqv_0(b_{l,i,m}) &= mv_p(\pi) + sv_0(b_{l,i,m}) + s(q-1)v_0(b_{l,i,m}) \\ &\geq w_s(b(\mathbf{w}_l)) + s(q-1)e. \end{aligned}$$

Suppose $\liminf_{l\to\infty} w_s(b(\mathbf{w}_l)) < L$ for some $L < \infty$. For l sufficiently large, we have $w_s(\lambda F^{-1}a(\mathbf{w}_l)) \geq L$ and $w_r(b(\mathbf{w}_l)) \geq L - h - w_s(\lambda^{-1}A)$; by the previous paragraph, this implies

$$\begin{aligned} w_s(b(\mathbf{w}_{l+1})) &\geq w_s(\mathbf{w}_{l+1}) \\ &= w_s(F\mathbf{v}_{l+1} - \lambda\mathbf{v}_{l+1}) \\ &= w_s(F\mathbf{v}_l + Ff(\mathbf{w}_l) - \lambda\mathbf{v}_l - \lambda f(\mathbf{w}_l)) \\ &= w_s(\mathbf{w}_l + Ff(\mathbf{w}_l) - \lambda f(\mathbf{w}_l)) \\ &= w_s(\lambda^{-1}Fb(\mathbf{w}_l) + \lambda F^{-1}a(\mathbf{w}_l)) \\ &\geq \min\{w_s(\lambda^{-1}A) + w_s(b(\mathbf{w}_l)) + s(q-1)e, L\}. \end{aligned}$$

We first deduce from this inequality that $w_s(b(\mathbf{w}_l))$ is bounded below: pick any l, choose $C < L$ such that $w_s(b(\mathbf{w}_l)) > C$, then note that $w_s(b(\mathbf{w}_{l+1})) \geq \min\{L, C + w_s(\lambda^{-1}A) + s(q-1)e\} > C$. If we put $M = \liminf w_s(b(\mathbf{w}_l))$, we thus have $-\infty < M < L$. However, in the inequality above, the limit inferior of the left side is M, while the limit inferior of the smaller right side is $\min\{L, M + w_s(\lambda^{-1}A) + s(q-1)e\} > M$. This contradiction shows that no L can exist as above, and so $w_s(\lambda^{-1}Fb(\mathbf{w}_l)) \to \infty$.

From $w_s(\lambda F^{-1}a(\mathbf{w}_l)) \to \infty$ for $0 < s \leq qr$, and $w_s(\lambda^{-1}Fb(\mathbf{w}_l)) \to \infty$ for $0 < s \leq r$, we conclude that $w_s(a(\mathbf{w}_l)) \to \infty$ and $w_s(b(\mathbf{w}_l)) \to \infty$ for $0 < s \leq r$. Thus $\mathbf{w}_l$ converges to 0 in the Fréchet topology, and $\mathbf{v}_l$ converges to a limit $\mathbf{v}$ satisfying $F\mathbf{v} = \lambda\mathbf{v}$.

Finally, we check that $\mathbf{v} \neq 0$. First note that $w_r(\lambda^{-1}[T]\mathbf{e}_1) = dr - v_p(\lambda)$, while

$$\begin{aligned} w_r(\mathbf{v}_0 - \lambda^{-1}[T]\mathbf{e}_1) &= w_r([T^{1/q}]F^{-1}\mathbf{e}_1) \\ &\geq dr/q + c \\ &> dr - v_p(\lambda) \end{aligned}$$

by our choice of d. Therefore $w_r(\mathbf{v}_0) = dr - v_p(\lambda)$. On the other hand,

$$\begin{aligned} w_r(\lambda^{-1}\mathbf{w}_0) &= w_r(\lambda^{-1}F\mathbf{v}_0 - \mathbf{v}_0) \\ &= w_r(\lambda^{-2}F[T]\mathbf{e}_1 + \lambda^{-1}[T]\mathbf{e}_1 - \lambda^{-1}[T]\mathbf{e}_1 - [T^{1/q}]F^{-1}\mathbf{e}_1) \\ &= w_r(\lambda^{-2}F[T]\mathbf{e}_1 - [T^{1/q}]F^{-1}\mathbf{e}_1) \\ &\geq \min\{rdq + c - 2v_p(\lambda), rd/q + c\}. \end{aligned}$$

We have just checked that the second term in braces is greater than $dr - v_p(\lambda) = w_r(\mathbf{v}_0)$. As for the first term,

$$rdq + c - 2v_p(\lambda) - (dr - v_p(\lambda)) = dr(q-1) + c - v_p(\lambda)$$

is positive, again by the choice of d. Therefore $w_r(\lambda^{-1}\mathbf{w}_0) > w_r(\mathbf{v}_0)$.

Since we showed earlier that $w_r(\mathbf{w}_l)$ is a strictly increasing function of l, we have $w_r(\lambda^{-1}\mathbf{w}_l) \geq w_r(\lambda^{-1}\mathbf{w}_0)$ for $l \geq 0$. We also showed earlier that $w_r(\mathbf{v}_{l+1} - \mathbf{v}_l) \geq w_r(\lambda^{-1}\mathbf{w}_l)$ for $l \geq 0$. Thus $w_r(\mathbf{v}_{l+1} - \mathbf{v}_l) \geq w_r(\lambda^{-1}\mathbf{w}_0)$ for each l, and so $w_r(\mathbf{v}_l - \mathbf{v}_0) \geq w_r(\lambda^{-1}\mathbf{w}_0)$. It follows that $w_r(\mathbf{v} - \mathbf{v}_0) \geq w_r(\lambda^{-1}\mathbf{w}_0) > w_r(\mathbf{v}_0)$; in particular, $\mathbf{v} \neq 0$, so λ and $\mathbf{v}$ satisfy the desired conditions. $\square$

Corollary 4.9. *Every σ-module M over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ admits a basis $\mathbf{v}_1, \dots, \mathbf{v}_n$ such that $\mathbf{v}_i$ is an eigenvector in $M/\operatorname{SatSpan}(\mathbf{v}_1, \dots, \mathbf{v}_{i-1})$ for $i = 1, \dots, n$.*

Proof. By the proposition and Lemma 4.1, every σ-module over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ contains a primitive eigenvector. The corollary now follows by induction on the rank of M. $\square$

Corollary 4.10. *The set of slopes of eigenvectors of M, over all finite extensions of $\mathcal{O}$, is bounded below.*

Proof. Combine the previous corollary with Proposition 4.4. $\square$

4.3 Raising the Newton polygon

In the previous section, we produced within any σ-module over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ a basis on which F acts by a triangular matrix. By Proposition 4.5, if there is a basis of eigenvectors, the valuations of the diagonal entries of this matrix majorize the slopes of the eigenvectors. Thus to produce a basis of eigenvectors, we need to "raise the Newton polygon", i.e., find eigenvectors whose eigenvalues have smaller slopes than the ones we started with. In this section, we carry this process out by direct computation in an important special case; the general process, using this case in some basic steps, will follow in the next section.

By a *Puiseux polynomial* over a field K, we shall mean a formal expression of the form

$$P(z) = \sum_{i \in I} c_i z^i$$

where I is a finite set of nonnegative rationals and $c_i \in K$ for each $i \in I$. If K has a valuation v_K, we define the Newton polygon of a Puiseux polynomial, by analogy with the definition for an ordinary polynomial, as the lower convex hull of the set of points $(-i, v_K(c_i))$. In fact, for some integer n, $P(z^n)$ is an ordinary polynomial; by comparing the Newton polygons of $P(z)$ and $P(z^n)$, and using the usual theory of Newton polygons of polynomials over fields complete with respect to a valuation, we obtain the following result.

Lemma 4.11. *Let $P(z)$ be a Puiseux polynomial over the t-adic completion of $k((t))^{\mathrm{alg}}$. Then P has a root of valuation l if and only if the Newton polygon of P has a segment of slope l.*

For $x \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ a strong semiunit, we refer to $v_0(x)$ as the *valuation* of x.

Lemma 4.12. *Let n be a positive integer, and let $x = \sum_{i=0}^n u_i \pi^i$ for some strong semiunits $u_i \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ of negative (or infinite) valuation, not all zero. Then the system of equations*

$$a^\sigma = \pi a, \qquad \pi b^{\sigma^n} = b - ax \tag{1}$$

has a solution with $a, b \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ not both zero.

Proof. For $i \in \{0, \dots, n\}$ for which $u_i \neq 0$, $l \in \mathbb{Z}$ and $m \in \mathbb{R}^+$, put

$$f(i,l,m) = (v_0(u_i) + mq^{-l})q^{-n(i+l)},$$

Note that for fixed i and m, $f(i,l,m)$ approaches 0 from below as $l \to +\infty$, and tends to $+\infty$ as $l \to -\infty$. Thus the minimum $h(m) = \min_{i,l}\{f(i,l,m)\}$ is well-defined. Observe that the map $h : \mathbb{R}^+ \to \mathbb{R}$ is continuous and piecewise linear with everywhere positive slope, and $h(qm) = q^{-n}h(m)$ because $f(i, l+1, qm) = q^{-n} f(i,l,m)$. Since $f(i,l,m)$ takes negative values for fixed i, l and small m, $h(m) < 0$ for some m, implying $h(q^j m) < 0$ for all $j \in \mathbb{Z}$, so h takes only negative values. We conclude that h is a continuous increasing bijection of $\mathbb{R}^+$ onto $\mathbb{R}^-$.

Pick $t \in \mathbb{R}^+$ at which h changes slope, let S be the finite set of ordered pairs (i,l) for which $f(i,l,t) < q^{-n}h(t)$, and let T be the set of ordered pairs (i,l) for which $f(i,l,t) < 0$; then T is infinite (and contains S), but the values of l for pairs $(i,l) \in T$ are bounded below. For each pair $(i,l) \in T$, put $s(i,l) = \lfloor \log_{q^n}(h(t)/f(i,l,t)) \rfloor$. This function has the following properties:

(a) $s(i,l) \geq 0$ for all $(i,l) \in T$;

(b) $f(i,l,t)q^{ns(i,l)} \in [h(t), q^{-n}h(t))$ for all $(i,l) \in T$;

(c) $(i,l) \in S$ if and only if $(i,l) \in T$ and $s(i,l) = 0$;

(d) for any $e > 0$, there are only finitely many pairs $(i,l) \in T$ such that $s(i,l) \leq e$.

For $c \in \mathbb{R}$, let U_c be the set of $z \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ such that $v_m(z) = \infty$ for $m < 0$ and $v_m(z) \geq c$ for $m \geq 0$. Then the function

$$r(z) = \sum_{(i,l)\in T} \pi^{s(i,l)} u_i^{\sigma^{-ni-nl+ns(i,l)}} z^{\sigma^{-ni-(n+1)l+ns(i,l)}}$$

is well-defined (by (d) above, the series is π-adically convergent) and carries U_t into $U_{h(t)}$ because for $z \in U_t$ and $m \geq 0$,

$$\begin{aligned} v_m\left(u_i^{\sigma^{-ni-nl+ns(i,l)}} z^{\sigma^{-ni-(n+1)l+ns(i,l)}}\right) &\geq q^{-ni-nl+ns(i,l)} v_0(u_i) + \min_j \{q^{-ni-(n+1)l+n(s,i,l)} v_j(z)\} \\ &\geq q^{ns(i,l)} q^{-n(i+l)} (v_0(u_i) + tq^{-l}) \\ &= q^{ns(i,l)} f(i,l,t) \\ &\geq h(t). \end{aligned}$$

The reduction of $r(z)$ modulo π is congruent to a finite sum over pairs $(i,l) \in S$, so it is a Puiseux polynomial in the reduction of z. Since $s(i,l) = 0$ for all $(i,l) \in S$ and the values $-ni-(n+1)l$ are all distinct (because i only runs over $\{0,\dots,n\}$), we get a distinct monomial modulo π for each pair $(i,l) \in S$.

We now consider the Newton polygon of the Puiseux polynomial given by the reduction of $r(z) - w$, for $w \in U_{h(t)}$. It is the convex hull of the points $(-q^{-ni-(n+1)l}, v_0(u_i)q^{-ni-nl})$ for each $(i,l) \in S$, together with $(0, v_0(w))$. The line $y = tx + h(t)$ either passes through or lies below the point corresponding to (i,l), depending on whether $f(i,l,t)$ is equal to or strictly greater than $h(t)$. Moreover, $(0, v_0(w))$ lies on or above the line because $v_0(w) \geq h(t)$. Since h changes slope at t, there must be at least two points on the line; therefore the Newton polygon has a segment of slope t. By Lemma 4.11, the Puiseux polynomial has a root of valuation t. In other words, there exists $z \in U_t$ with $v_0(z) = t$ such that $r(z) \equiv w \pmod{\pi}$.

As a consequence of the above reasoning, we see that the image of U_t is dense in $U_{f(t)}$ with respect to the π-adic topology. Since U_t is complete, U_t must surject onto $U_{f(t)}$. Moreover, we can take $w = 0$ and obtain $z_0 \in U_t$ with $v_0(z_0) = t$ such that $r(z_0) \equiv 0 \pmod{\pi}$; in particular, z_0 is nonzero modulo π. We may then obtain $z_1 \in U_t$ such that $r(z_1) = r(z_0)/\pi$. Put $z = z_0 - \pi z_1$; then $z \not\equiv 0 \pmod{\pi}$ and so is nonzero, but $r(z) = 0$.

Now set $a = \sum_{l=-\infty}^{\infty} \pi^l z^{\sigma^{-l}}$; the sum converges in $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ because for $s > 0$, $w_s(\pi^l z^{\sigma^{-l}}) \geq l v_p(\pi) + rq^{-l}t$ and the latter tends to ∞ as $l \to \pm\infty$ (because $t > 0$). Then

$$\begin{aligned} ax &= \sum_{i=0}^{n} \sum_{l=-\infty}^{\infty} \pi^{i+l} u_i z^{\sigma^{-l}} \\ &= \sum_{(i,l)\in T} \pi^{i+l} u_i z^{\sigma^{-l}} + \sum_{(i,l)\notin T} \pi^{i+l} u_i z^{\sigma^{-l}}. \end{aligned}$$

Let A and B denote the two sums in the last line; then $v_m(B) \geq 0$ for all m, so by Proposition 3.19(d) (with σ replaced by σ^n), B can be written as $\pi b_1^{\sigma^n} - b_1$ for some $b_1 \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$. On the other hand, we claim that A can be rewritten as $r(z) + \pi b_2^{\sigma^n} - b_2$ for

$$\begin{aligned} b_2 &= \sum_{(i,l)\in T} \sum_{j=1}^{i+l-s(i,l)} \pi^{i+l-j} u_i^{\sigma^{-nj}} z^{\sigma^{-l-nj}} \\ &= \sum_{(i,l)\in T} \sum_{k=0}^{i+l-s(i,l)-1} \pi^{k+s(i,l)} \left(u_i^{\sigma^{-ni-nl+ns(i,l)}} z^{\sigma^{-ni-(n+1)l+ns(i,l)}}\right)^{\sigma^{nk}} \end{aligned}$$

(via the substitution $k = i + l - s(i,l) - j$); we must check that this series converges π-adically and that its limit is overconvergent. Note that as $l \to +\infty$ for i fixed, $f(i,l,m)$ is asymptotic to $v_0(u_i)q^{-n(i+l)}$. Therefore $i + l - s(i,l)$ is bounded, so the possible values of k are uniformly bounded over all pairs $(i,l) \in T$. This implies on one hand that the series converges π-adically (since l is bounded below over pairs $(i,l) \in T$ and $s(i,l) \to \infty$ as $l \to \infty$), and on the other hand that $v_m(b_2)$ is bounded below uniformly in m (since the quantity in parentheses in the second sum belongs to $U_{h(t)}$), so $b_2 \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$.

Having shown that the series defining b_2 converges, we can now verify that $b_2 - \pi b_2^{\sigma^n} = r(z) - A$: the quantity on the left is the sum over pairs $(i,l) \in T$ of a sum over k which telescopes, leaving the term $k = 0$ minus the term $k = i + l - s(i,l)$, or

$$\pi^{s(i,l)} u_i^{\sigma^{-ni-nl+ns(i,l)}} z^{\sigma^{-ni-(n+1)l+ns(i,l)}} - \pi^{i+l} u_i z^{\sigma^{-l}},$$

which when summed over pairs $(i,l) \in T$ yields $r(z) - A$.

Since $r(z) = 0$ by construction, we have $\pi b^{\sigma^n} = b - ax$ for $b = -(b_1 + b_2)$. Thus (a,b) constitute a solution of (1), as desired. □

We apply the previous construction to study the system of equations

$$a^\sigma = \pi a, \qquad \pi b^{\sigma^n} = b - ac, \tag{2}$$

where $c \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ is given. Notice that replacing c by $c + \pi^{n+1} y^{\sigma^n} - y$ does not alter whether (2) has a solution: for any a such that $a^\sigma = \pi a$, if $\pi b^{\sigma^n} = b - ac$, then

$$\pi(b - ay)^{\sigma^n} = (b - ay) - a(c + \pi^{n+1} y^{\sigma^n} - y).$$

We begin by analyzing (2) in a restricted case.

Lemma 4.13. *For any positive integer n and any $c \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ such that $v_m(c) \geq -1$ for all m and $v_m(c) = \infty$ for some m, there exist $a, b \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ not both zero satisfying* (2).

Proof. By multiplying c by a power of π, we may reduce to the case where $v_m(c) = \infty$ for $m < 0$. Define the sequences $c_0, c_1, \dots$ and $d_0, d_1, \dots$ as far as is possible by the following iteration. First put $c_0 = c$ and $d_0 = 0$. Given c_i, if $v_0(c_i) < -1/q^n$, stop. Otherwise, let d_i be a strong semiunit congruent to c_i modulo π and put $c_{i+1} = (c_i + \pi^{n+1} d_i^{\sigma^n} - d_i)/\pi$. Note that $v_m(c_i) \geq -1$ and $v_m(d_i) \geq -1/q^n$ for all $m \geq 0$ and all i.

If the iteration never terminates, then we have $c + \pi^{n+1} d^{\sigma^n} - d = 0$ for $d = \sum_{i=0}^\infty d_i \pi^i$. In this case, apply Proposition 3.19(b) to produce a nonzero such that $a^\sigma = \pi a$ and set $b = ad$ to obtain a solution to (2).

If the iteration terminates at c_l, set $d = \sum_{i=0}^{l-1} d_i \pi^i$, so that $\pi^l c_l = c + \pi^{n+1} d^{\sigma^n} - d$. Let $\sum_{j=0}^\infty u_j \pi^j$ be a strong semiunit decomposition of c_l, necessarily having $v_0(u_0) < -1/q^n$. Put $e = \sum_{j=n+1}^\infty u_j^{\sigma^{-n}} \pi^{j-n-1}$ and set $x = c_l - \pi^{n+1} e^{\sigma^n} + e$. Then

$$x = \sum_{j=1}^n \pi^j u_j + \left(u_0 + \sum_{j=n+1}^\infty u_j^{\sigma^{-n}} \pi^{j-n-1} \right),$$

and the quantity in parentheses is a strong semiunit of the same valuation as u_0, since $v_0(u_0) < -1/q^n \leq v_0(u_j^{\sigma^{-n}})$ for all j. Thus x satisfies the condition of Lemma 4.12, so there exist $a', b' \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ not both zero so that

$$(a')^\sigma = \pi a', \qquad \pi (b')^{\sigma^n} = b' - a'x.$$

We obtain a solution of (2) by setting $a = a'$, $b = a'd - \pi^l a'e + \pi^l b'$. ☐

We now analyze (2) in general by reducing to the special case treated above.

Lemma 4.14. *For any positive integer n and any $c \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$, there exist $a, b \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ not both zero such that* (2) *holds.*

Proof. Let $\sum_i u_i \pi^i$ be a strong semiunit decomposition of c, and let N be the smallest integer such that $v_0(u_N) < 0$, or ∞ if there is no such integer. By Proposition 3.19(d), there exists $y \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ such that $\pi^{n+1} y^{\sigma^n} - y + \sum_{i=-\infty}^{N-1} u_i \pi^i = 0$.

If $N = \infty$, then in fact $\pi^{n+1} y^{\sigma^n} - y + c = 0$, so we obtain a solution of (2) by choosing a nonzero with $a^\sigma = \pi a$ via Proposition 3.19(b), and setting $b = ay$. So suppose hereafter that $N < \infty$.

For each $i \geq N$ for which $u_i \neq 0$, set $t_i = \lceil \log_{q^n}(-v_0(u_i)) \rceil$, so that $-1 \leq v_0(u_i^{\sigma^{-nt_i}}) < -1/q^n$ for all such i. Then the sum

$$z = \sum_{i=N}^{\infty} \sum_{j=1}^{t_i} u_i^{\sigma^{-nj}} \pi^{i-(n+1)j}$$

is π-adically convergent: $-v_0(u_i)$ grows at most linearly in i, so t_i grows at most logarithmically and $i - (n+1)t_i \to \infty$ as $i \to \infty$. Moreover, t_i is bounded below, so $v_0(u_i^{\sigma^{-nj}} \pi^{i-(n+1)j})$ is as well; thus the sum z is in $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$.

Put $c' = c + \pi^{n+1}(y-z)^{\sigma^n} - (y-z)$; then

$$\begin{aligned} c' &= \sum_{i=N}^{\infty} \left(u_i \pi^i + \sum_{j=1}^{t_i} u_i^{\sigma^{-nj}} \pi^{i-(n+1)j} - \sum_{j=1}^{t_i} u_i^{\sigma^{-n(j-1)}} \pi^{i-(n+1)(j-1)} \right) \\ &= \sum_{i=N}^{\infty} u_i^{\sigma^{-nt_i}} \pi^{i-(n+1)t_i}, \end{aligned}$$

so that $v_m(c') \geq -1$ for all m. By Lemma 4.13, there exist $a', b' \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ not both zero such that

$$(a')^\sigma = \pi a', \qquad \pi (b')^{\sigma^n} = b' - a'c';$$

we obtain a solution of (2) by setting $a = a', b = b' + a'(y-z)$. ☐

We now prove our basic result on raising the Newton polygon, i.e., reducing the slope of an eigenvector.

Proposition 4.15. *Let m and n be positive integers, and let M be a σ-module over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ admitting a basis $\mathbf{v}_1, \dots, \mathbf{v}_n, \mathbf{w}$ such that for some $c_i \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$,*

$$\begin{aligned} F\mathbf{v}_i &= \mathbf{v}_{i+1} \qquad (i = 1, \dots, n-1) \\ F\mathbf{v}_n &= \pi \mathbf{v}_1 \\ F\mathbf{w} &= \pi^{-m}\mathbf{w} + c_1\mathbf{v}_1 + \cdots + c_n \mathbf{v}_n. \end{aligned}$$

Then there exists $\mathbf{y} \in M$ such that $F\mathbf{y} = \mathbf{y}$.

This will ultimately be a special case of our main results; what makes this case directly tractable is that if $\mathrm{SatSpan}(\mathbf{v}_1, \dots, \mathbf{v}_n)$ does not admit an F-stable complement in M (i.e., is not a direct summand of M in the category of σ-modules), then the map $\mathbf{y} \mapsto F\mathbf{y} - \mathbf{y}$ is actually surjective, as predicted by the expected behavior of the special Newton polygon.

Proof. Suppose $\mathbf{y} = d\mathbf{w} + b_1\mathbf{v}_1 + \cdots + b_n\mathbf{v}_n$ satisfies $F\mathbf{y} = \mathbf{y}$, or in other words

$$d\mathbf{w} + \sum_{i=1}^n b_i \mathbf{v}_i = \pi^{-m} d^\sigma \mathbf{w} + \sum_{i=1}^n d^\sigma c_i \mathbf{v}_i + \sum_{i=1}^{n-1} b_i^\sigma \mathbf{v}_{i+1} + \pi b_n^\sigma \mathbf{v}_1.$$

Comparing coefficients in this equation, we have $b_i^\sigma = b_{i+1} - d^\sigma c_{i+1}$ for $i = 1, \dots, n-1$, as well as $\pi b_n^\sigma = b_1 - d^\sigma c_1$ and $d^\sigma = \pi^m d$. If we use the first n relations to eliminate $b_2, \dots, b_n$, we get

$$\begin{aligned} b_1^{\sigma^n} &= b_2^{\sigma^{n-1}} - d^{\sigma^n} c_2^{\sigma^{n-1}} \\ &= b_3^{\sigma^{n-2}} - d^{\sigma^{n-1}} c_3^{\sigma^{n-2}} - d^{\sigma^n} c_2^{\sigma^{n-1}} \\ &\;\;\vdots \\ &= b_n^\sigma - d^{\sigma^2} c_n^\sigma - \cdots - d^{\sigma^n} c_2^{\sigma^{n-1}} \\ &= \pi^{-1} b_1 - d(\pi^{m-1} c_1 + \pi^{2m} c_n^\sigma + \pi^{3m} c_{n-1}^{\sigma^2} + \cdots + \pi^{nm} c_2^{\sigma^{n-1}}). \end{aligned}$$

Let c' be the quantity in parentheses in the last line. We have shown that if $F\mathbf{y} = \mathbf{y}$ has a nonzero solution, then the system of equations

$$d^\sigma = \pi^m d, \qquad \pi b_1^{\sigma^n} = b_1 - \pi c' d \tag{3}$$

has a solution with b_1, d not both zero. Conversely, from any nonzero solution of (3) we may construct a nonzero $\mathbf{y} \in M$ such that $F\mathbf{y} = \mathbf{y}$, by using the relations $b_i^\sigma = b_{i+1} - d^\sigma c_{i+1}$ to successively define $b_2, \dots, b_n$.

By Proposition 3.19(b), we can find $e \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ nonzero such that $e^\sigma = \pi^{m-1} e$; we will construct a solution of (3) with $d = ae$ for some a such that $a^\sigma = \pi a$. Namely, put $c = \pi c' e$, and apply Lemma 4.14 to find $a, b \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$, not both zero, such that

$$a^\sigma = \pi a, \qquad \pi b^{\sigma^n} = b - ac.$$

Then $b_1 = b$ and $d = ae$ constitute a nonzero solution of (2); as noted above, this implies that there exists $\mathbf{y} \in M$ nonzero with $F\mathbf{y} = \mathbf{y}$, as desired. □

4.4 Construction of the special Newton polygon

We now assemble the results of the previous sections into the following theorem, the main result of this chapter.

Theorem 4.16. *Let M be a σ-module over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$. Then M can be expressed as a direct sum of standard σ-submodules.*

As the proof of this theorem is somewhat intricate, we break off parts of the argument into separate lemmas. In these lemmas, a "suitable extension" of $\mathcal{O}[\frac{1}{p}]$ means one whose value group contains whatever slope is desired to be the slope of an eigenvector. By Proposition 4.7, proving the existence of an eigenvector of prescribed slope over a single suitable extension implies the same over any suitable extension.

Lemma 4.17. *Let M be a σ-module over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ of rank 1, and suppose F acts on some generator $\mathbf{v}$ via $F\mathbf{v} = c\mathbf{v}$. Then M contains an eigenvector, and any eigenvector has slope $v_p(c)$.*

Note that $v_n(c) = \infty$ for some n by Corollary 3.23, so that $v_p(c)$ makes sense.

Proof. The existence of an eigenvector of slope $v_p(c)$ follows from Proposition 3.18. The uniqueness of the slope follows from Corollary 4.6. ☐

For M of rank 1, we call this unique slope the slope of M. Note that if $0 \to L \to M \to N \to 0$ is an exact sequence of σ-modules and L, M, N have ranks l, m, n, respectively, then the slope of $\wedge^m M$ is the sum of the slopes of $\wedge^l L$ and $\wedge^n N$. (This assertion will be vastly generalized by Proposition 5.13 later.)

Lemma 4.18. *Let M be a σ-module over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ of rank 2, and let d be the slope of $\wedge^2 M$. Then M contains an eigenvector of slope $d/2$ over a suitable extension of $\mathcal{O}[\frac{1}{p}]$.*

Proof. We may assume without loss of generality that $d/2$ belongs to the value group of $\mathcal{O}[\frac{1}{p}]$. Let e be the smallest integer such that M contains an eigenvector of slope $ev_p(\pi)$. (There is such an integer by Proposition 4.8, and there is a smallest one by Corollary 4.10.) By twisting, we may reduce to the case where $e = 1$.

Put $m = 1 - (d/v_p(\pi))$ and suppose by way of contradiction that $m > 0$. Choose a eigenvector $\mathbf{v}$ with $F\mathbf{v} = \pi\mathbf{v}$, which is necessarily primitive by Lemma 4.1; then by Lemma 4.17 applied to $M/\operatorname{SatSpan}(\mathbf{v})$, we can find $\mathbf{w}$ such that $\mathbf{v}, \mathbf{w}$ form a basis of M and $F\mathbf{w} = \pi^{-m}\mathbf{w} + c\mathbf{v}$ for some $c \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$. Now by Proposition 4.15, M contains an eigenvector $\mathbf{v}_1$ with $F\mathbf{v}_1 = \mathbf{v}_1$, contradicting the definition of e.

Hence $m \leq 0$, which implies $d \geq v_p(\pi)$. Since $d/2$ is also a multiple of $v_p(\pi)$, we must have $d/2 \geq v_p(\pi)$; by Proposition 4.2, M contains an eigenvector of slope $d/2$. ☐

Lemma 4.19. *Let M be a σ-module over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ of rank n, and let d be the slope of $\wedge^n M$. Then M contains eigenvectors of all slopes greater than d/n over suitable extensions of $\mathcal{O}[\frac{1}{p}]$.*

Proof. We proceed by induction on n. The case $n = 1$ follows from Lemma 4.17, and the case $n = 2$ follows from Lemma 4.18. Suppose $n > 2$ and that the lemma has been proved for all smaller values of n. Let s be the greatest lower bound of the set of rational numbers that occur as slopes of eigenvectors of M (over suitable extensions of $\mathcal{O}[\frac{1}{p}]$). Again, the set is nonempty by Proposition 4.8 and is bounded below by Corollary 4.10.

For each $\epsilon > 0$ such that $s + \epsilon \in \mathbb{Q}$, over a suitable extension of $\mathcal{O}[\frac{1}{p}]$ there exist an eigenvector $\mathbf{v}$ of M of slope $s + \epsilon$ and (by the induction hypothesis) an eigenvector $\mathbf{w}$ of $M/\operatorname{SatSpan}(\mathbf{v})$ of slope at most $s' = (d - s - \epsilon)/(n - 1) + \epsilon$. The preimage of $\operatorname{SatSpan}(\mathbf{w})$ in M has rank 2, so is covered by the induction hypothesis; it thus contains, for any $\delta > 0$, an eigenvector of slope at most

$$\frac{s+\epsilon}{2} + \frac{d - s + (n-2)\epsilon}{2(n-1)} + \delta$$

over a suitable extension of $\mathcal{O}[\frac{1}{p}]$. Such an eigenvector is also an eigenvector of M, so its slope is at least s. Letting ϵ and δ go to 0 in the resulting inequality yields

$$\frac{s}{2} + \frac{d-s}{2(n-1)} \geq s,$$

which simplifies to $s \leq d/n$, as desired. $\square$

Lemma 4.20. *Let M be a σ-module over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ of rank n, and let d be the slope of $\wedge^n M$. Then M contains an eigenvector of slope d/n over a suitable extension of $\mathcal{O}[\frac{1}{p}]$.*

Proof. We proceed by induction on n; again, the case $n = 1$ follows from Lemma 4.17 and the case $n = 2$ follows from Lemma 4.18. Without loss of generality, we may assume the value group of $\mathcal{O}$ contains d/n, and then that $d = 0$.

By Lemma 4.19, there exists an eigenvector $\mathbf{v}$ of M of slope $v_p(\pi)/(n-1)$ over $\mathcal{O}[\pi^{1/(n-1)}]$; we may as well assume $F\mathbf{v} = \pi^{1/(n-1)}\mathbf{v}$. Let N be the saturated span of $\mathbf{v}$ and its conjugates over $\mathcal{O}[\frac{1}{p}]$; let m be the rank of N and s the slope of $\wedge^m N$. Then $m \leq n - 1$ and $s \leq mv_p(\pi)/(n-1)$. If $m < n-1$, then $0 < mv_p(\pi)/(n-1) < v_p(\pi)$, so $s \leq 0$ and the induction hypothesis implies that N contains an eigenvector of slope 0. The same argument applies if $m = n - 1$ and $s < v_p(\pi)$.

Suppose instead that $m = n - 1$ and $s = v_p(\pi)$. Write $\mathbf{v} = \mathbf{v}_1 + \pi^{-1/(n-1)}\mathbf{v}_2 + \cdots + \pi^{-(n-2)/(n-1)}\mathbf{v}_{n-1}$ with each $\mathbf{v}_i$ defined over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ (with no extension of $\mathcal{O}[\frac{1}{p}]$); then $\mathbf{v}_1, \dots, \mathbf{v}_{n-1}$ are linearly independent in N, and we have $F\mathbf{v}_i = \mathbf{v}_{i+1}$ for $i = 1, \dots, n-1$ and $F\mathbf{v}_{n-1} = \pi\mathbf{v}_1$. In particular, the $\mathbf{v}_i$ must be a basis of N or else $\wedge^{n-1}N$ would have slope less than s. The slope of M/N is $-v_p(\pi)$, so by Lemma 4.17, we can choose $\mathbf{w} \in M$ such that $F\mathbf{w} \equiv \pi^{-1}\mathbf{w} \pmod{N}$. Proposition 4.15 then implies that M contains an eigenvector of slope 0, as desired. $\square$

Proof of Theorem 4.16. We proceed by induction on the rank of M. If $\operatorname{rank} M = 1$, then M is standard by Lemma 4.17. Suppose $\operatorname{rank} M = n > 1$, and that the proposition has been established for all σ-modules of rank less than n. For any rational number c, define the *$\mathcal{O}$-index* of c as the smallest integer m such that mc lies in the value group of $\mathcal{O}[\frac{1}{p}]$. The set of rational numbers of $\mathcal{O}$-index less than or equal to n which occur as slopes of eigenvectors of M is discrete (obvious), nonempty (by Proposition 4.8), and bounded below (by Corollary 4.10), so has a smallest element r.

Let d be the slope of $\wedge^n M$. By Lemma 4.20, we have $r \leq d/n$. Let s be the $\mathcal{O}$-index of r, and let λ be an element of a degree s extension $\mathcal{O}'[\frac{1}{p}]$ of $\mathcal{O}[\frac{1}{p}]$ such that $v_p(\lambda) = r$ and $\lambda^s \in \mathcal{O}[\frac{1}{p}]$. Choose an eigenvector $\mathbf{v}$ over $\mathcal{O}'[\frac{1}{p}]$ with $F\mathbf{v} = \lambda\mathbf{v}$, and write $\mathbf{v} = \sum_{i=0}^{s-1} \lambda^{-i}\mathbf{w}_i$ for $\mathbf{w}_i \in M$, so that $F\mathbf{w}_i = \mathbf{w}_{i+1}$ for $i = 0, \dots, s-2$ and $F\mathbf{w}_{s-1} = \lambda^s\mathbf{w}_0$. Put $N = \mathrm{SatSpan}(\mathbf{w}_0, \dots, \mathbf{w}_{s-1})$ and $m = \operatorname{rank} N$; then $s \geq m$, and the slope of $\wedge^m N$ is at most mr, since N is the saturated span of eigenvectors of slope r.

If $m = n$, then also $s = n$ and $\mathbf{w}_0 \wedge \cdots \wedge \mathbf{w}_{n-1}$ is an eigenvector of $\wedge^n M$ of slope rn. Thus $rn \geq d$; since $r \leq d/n$ as shown earlier, we conclude $r = d/n$, $\mathbf{w}_0, \dots, \mathbf{w}_{n-1}$ form a basis of M, and M is standard, completing the proof in this case. Thus we assume $m < n$ hereafter.

Given that $m < n$, we may apply the induction hypothesis to N, deducing in particular that its smallest slope is at most r and has $\mathcal{O}$-index not greater than m. This yields a contradiction unless that slope is r, which is only possible if the slope of $\wedge^m N$ is mr. In turn, mr belongs to the value group of $\mathcal{O}[\frac{1}{p}]$ only if $m = s$. Thus $m = s$, and since $\mathbf{w}_0 \wedge \cdots \wedge \mathbf{w}_{s-1}$ is an eigenvector of N of slope rs, $\mathbf{w}_0, \dots, \mathbf{w}_{s-1}$ form a basis of N, and N is standard.

Apply the induction hypothesis to M/N to express it as a sum $P_1 \oplus \cdots \oplus P_l$ of standard σ-submodules; note that the $\mathcal{O}$-index of the slope of P_i divides the rank of P_i, and so is at most n. If $l = 1$, then the slope of P_1 cannot be less than r (else the slope of $\wedge^n M$ would be less than d), so by Proposition 4.3, M can be split as a direct sum of N with a standard σ-module. If $l > 1$, let M_i be the preimage of P_i under the projection $M \to M/N$; again the slope of each P_i cannot be less than r, else the induction hypothesis would imply that M_i contains an eigenvector of slope less than r and $\mathcal{O}$-index not exceeding n, contradiction. Thus by Proposition 4.3 again, each M_i can be split as a direct sum $N \oplus N_i$ of σ-submodules, and we may decompose M as $N \oplus N_1 \oplus \cdots \oplus N_l$. This completes the induction in all cases. □

By Corollary 4.6, the multiset union of the slopes of the standard summands of a σ-module M over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ (each summand contributing its slope as many times as its rank) does not depend on the decomposition. Thus we define the *special Newton polygon* of M as the polygon with vertices (i, y_i) $(i = 0, \dots, n)$, where $y_0 = 0$ and $y_i - y_{i-1}$ is the i-th smallest slope of M (counting multiplicity). We extend this definition to σ-modules over $\Gamma_{\mathrm{an,con}}$ by base extending to $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$.

5 The generic Newton polygon

In this chapter, we recall the construction of the generic Newton polygon associated to a σ-module over Γ. The construction uses a classification result, the Dieudonné-Manin classification, for σ-modules over a complete discrete valuation ring with algebraically closed residue field. This classification does not descend very well, so we describe some weaker versions of the classification that can be accomplished under less restrictive conditions. These weaker versions either appear in or are inspired directly by [dJ].

5.1 Properties of eigenvectors

Throughout this section, let R be a discrete valuation ring which is unramified over $\mathcal{O}$. Again, we call an element $\mathbf{v}$ of a σ-module M over R or $R[\frac{1}{p}]$ an *eigenvector* if there exists $\lambda \in \mathcal{O}$ or $\mathcal{O}[\frac{1}{p}]$, respectively, such that $F\mathbf{v} = \lambda\mathbf{v}$, and refer to $v_p(\lambda)$ as the *slope* of $\mathbf{v}$. We call an eigenvector *primitive* if it forms part of a basis of M, but this definition is not very useful: every eigenvector is a $\mathcal{O}$-multiple of a primitive eigenvector of the same slope. In fact, in contrast with the situation over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$, the slopes of eigenvectors over R are "rigid".

Proposition 5.1. *Let M be a σ-module over $R[\frac{1}{p}]$, with k algebraically closed. Suppose M admits a basis $\mathbf{v}_1, \dots, \mathbf{v}_n$ of eigenvectors. Then any eigenvector $\mathbf{w}$ is an $\mathcal{O}[\frac{1}{p}]$-linear combination of those $\mathbf{v}_i$ of the same slope. In particular, any eigenvector has the same slope as one of the $\mathbf{v}_i$.*

Proof. Suppose $F\mathbf{v}_i = \lambda_i \mathbf{v}_i$ for some $\lambda_i \in \mathcal{O}[\frac{1}{p}]$, and write $\mathbf{w} = \sum_i c_i \mathbf{v}_i$ with $c_i \in R[\frac{1}{p}]$. If $F\mathbf{w} = \mu\mathbf{w}$ for $\mu \in \mathcal{O}[\frac{1}{p}]$, then equating the coefficients of $\mathbf{v}_i$ yields $\lambda_i c_i^\sigma = \mu c_i$. If $v_p(\lambda_i) \neq v_p(\mu)$, this forces $c_i = 0$; if $v_p(\lambda_i) = v_p(\mu)$, it forces $c_i \in \mathcal{O}[\frac{1}{p}]$. This proves the claim. $\square$

By imitating the proof of Proposition 4.5 using Proposition 5.1 in lieu of Proposition 4.4, we obtain the following analogue of Corollary 4.6.

Proposition 5.2. *Let M be a σ-module over $R[\frac{1}{p}]$. Suppose $\mathbf{v}_1, \dots, \mathbf{v}_n$ and $\mathbf{w}_1, \dots, \mathbf{w}_n$ are bases of eigenvectors with $F\mathbf{v}_i = \lambda_i \mathbf{v}_i$ and $F\mathbf{w}_i = \mu_i \mathbf{w}_i$, for some $\lambda_i, \mu_i \in \mathcal{O}[\frac{1}{p}]$. Then the sequences $v_p(\lambda_1), \dots, v_p(\lambda_n)$ and $v_p(\mu_1), \dots, v_p(\mu_n)$ are permutations of each other.*

In case M has a full set of eigenvectors of one slope, we have the following decomposition result.

Proposition 5.3. *Suppose k is algebraically closed, and let M be a σ-module over R spanned by eigenvectors of a single slope over $R \otimes_{\mathcal{O}} \mathcal{O}'$, for some finite extension $\mathcal{O}'$ of $\mathcal{O}$. Then M is isogenous to the direct sum of standard σ-modules of that slope.*

Proof. Let s be the common slope, and let m be the smallest positive integer such that ms is a multiple of $v_p(\pi)$. Since k is algebraically closed, there exists $\lambda \in \mathcal{O}'$ such that $\lambda^m \in \mathcal{O}$. Let $\mathcal{O}''$ be the integral closure of $\mathcal{O}$ in $\mathcal{O}[\frac{1}{p}](\lambda)$.

Note that M is spanned over $R \otimes_{\mathcal{O}} \mathcal{O}'$ by eigenvectors $\mathbf{v}$ with $F\mathbf{v} = \lambda\mathbf{v}$ because k is algebraically closed: if $F\mathbf{w} = \mu\mathbf{w}$ for some μ with $v_p(\mu) = v_p(\lambda)$, we can find $c \in \mathcal{O}'$ nonzero such that $c^\sigma = (\lambda/\mu)c$ and obtain a new eigenvector $\mathbf{v} = c\mathbf{w}$ with $F\mathbf{v} = \lambda\mathbf{v}$. We next verify that M is also spanned over $R \otimes_{\mathcal{O}} \mathcal{O}''$ by eigenvectors $\mathbf{v}$ with $F\mathbf{v} = \lambda\mathbf{v}$. Let $\mu_1, \dots, \mu_n$ be a basis of $\mathcal{O}'$ over $\mathcal{O}''$ consisting of elements fixed by σ (possible because k is algebraically closed). If $\mathbf{v}$ is an eigenvector over $R \otimes_{\mathcal{O}} \mathcal{O}'$ with $F\mathbf{v} = \lambda\mathbf{v}$, we can write $\mathbf{v} = \sum_i \mu_i \mathbf{w}_i$ for some $\mathbf{w}_i$ over $R \otimes_{\mathcal{O}} \mathcal{O}''$, and we must have $F\mathbf{w}_i = \lambda\mathbf{w}_i$ for each i. Thus $\mathbf{v}$ is in the span of the $\mathbf{w}_i$, so the span of eigenvectors of eigenvalue λ over $R \otimes_{\mathcal{O}} \mathcal{O}''$ has full rank over $R \otimes_{\mathcal{O}} \mathcal{O}'$, and thus has full rank over $R \otimes_{\mathcal{O}} \mathcal{O}''$.

Finally, we establish that M is isogenous to a direct sum of standard σ-modules. Let $\mathbf{v}$ be an eigenvector of eigenvalue λ over $R \otimes_{\mathcal{O}} \mathcal{O}''$; we can write $\mathbf{v} = \sum_{i=0}^{m-1} \mathbf{w}_i \lambda^{-i}$ for some $\mathbf{w}_i \in M$. Then $F\mathbf{w}_i = \mathbf{w}_{i+1}$ for $i = 0, \dots, m-2$ and $F\mathbf{w}_{m-1} = \lambda^m \mathbf{w}_0$, so the span of $\mathbf{w}_0, \dots, \mathbf{w}_{m-1}$ is standard. (Notice that $\mathbf{w}_0, \dots, \mathbf{w}_{m-1}$ must be linearly independent, otherwise any eigenvalue of their span would lie in an extension of $\mathcal{O}$ of degree strictly less than m, contrary to Proposition 5.2.) Let M_1 be the standard submodule just produced. Next, choose an eigenvector of eigenvalue λ linearly independent from M_1, and produce another standard submodule M_2. Then choose an eigenvector linearly independent from $M_1 \oplus M_2$, and so on until M is exhausted. □

5.2 The Dieudonné-Manin classification

Again, let R be a discrete valuation ring unramified over $\mathcal{O}$.

Lemma 5.4. *Suppose that R is complete with algebraically closed residue field. Given elements $a_0, \dots, a_{n-1}$ of R with a_0 nonzero, let M be the σ-module with basis $\mathbf{v}_1, \dots, \mathbf{v}_n$ such that*

$$\begin{aligned} F\mathbf{v}_i &= \mathbf{v}_{i+1} \qquad (i = 1, \dots, n-1) \\ F\mathbf{v}_n &= a_0\mathbf{v}_1 + \cdots + a_{n-1}\mathbf{v}_n. \end{aligned}$$

Suppose s belongs to the value group of R. Then the maximum number of linearly independent eigenvectors of slope s in M is less than or equal to the multiplicity m of s as a slope of the Newton polygon of the polynomial $x^n + a_{n-1}x^{n-1} + \cdots + a_0$ over R. Moreover, if $m > 0$, then M admits an eigenvector of slope s.

Proof. Let $l = \min_j\{-js + v_p(a_{n-j})\}$ (setting $a_n = 1$ for consistency); then there exists an index i such that $l = -js + v_p(a_{n-j})$ for $j = i$, $j = i + m$, and possibly for some values of $j \in \{i+1, \dots, i+m-1\}$, but not for any other values.

Let λ be an element of valuation s fixed by σ. Suppose $\mathbf{w} = \sum_j c_j \mathbf{v}_j$ satisfies $F\mathbf{w} = \lambda\mathbf{w}$. Then $\lambda c_1 = a_0 c_n^\sigma$ and $\lambda c_j = a_{j-1}c_n^\sigma + c_{j-1}^\sigma$ for $j = 2, \dots, n$. Solving for c_n yields

$$\begin{aligned} c_n &= \lambda^{-1}a_{n-1}c_n^\sigma + \lambda^{-1}c_{n-1}^\sigma \\ &= \lambda^{-1}a_{n-1}c_n^\sigma + \lambda^{-2}a_{n-2}^\sigma c_n^{\sigma^2} + \lambda^{-2}c_{n-2}^{\sigma^2} \\ &\;\;\vdots \\ &= \lambda^{-1}a_{n-1}c_n^\sigma + \lambda^{-2}a_{n-2}^\sigma c_n^{\sigma^2} + \cdots + \lambda^{-n+1}a_1^{\sigma^{n-2}} c_n^{\sigma^{n-1}} + \lambda^{-n+1}c_1^{\sigma^{n-1}} \\ &= \lambda^{-1}a_{n-1}c_n^\sigma + \lambda^{-2}a_{n-2}^\sigma c_n^{\sigma^2} + \cdots + \lambda^{-n+1}a_1^{\sigma^{n-2}} c_n^{\sigma^{n-1}} + \lambda^{-n}a_0^{\sigma^{n-1}} c_n^{\sigma^n}. \end{aligned}$$

In other words, $f(c_n) = 0$, where

$$f(x) = -x + \frac{a_{n-1}}{\lambda} x^{\sigma} + \frac{a_{n-2}^{\sigma}}{\lambda^2} x^{\sigma^2} + \cdots + \frac{a_0^{\sigma^{n-1}}}{\lambda^n} x^{\sigma^n}.$$

The coefficients of f of minimal valuation are on x^{σ^i}, $x^{\sigma^{i+m}}$, and possibly some in between.

Now suppose $\mathbf{w}_1, \dots, \mathbf{w}_{m+1}$ are linearly independent eigenvectors of M with $F\mathbf{w}_h = \lambda \mathbf{w}_h$ for $h = 1, \dots, m+1$. Write $\mathbf{w}_h = \sum_j c_{hj} \mathbf{v}_j$. Then $c_{1n}, \dots, c_{(m+1)n}$ are linearly independent over $\mathcal{O}_0$: if there were a relation $\sum_h d_h c_{hn} = 0$ with $d_h \in \mathcal{O}_0$ not all zero, we would have

$$\lambda \sum_h d_h c_{hj} = \left(\sum_h d_h c_{h(j-1)} \right)^{\sigma} + a_{j-1} \left(\sum_h d_h c_{hn} \right)^{\sigma} \qquad (j = 2, \dots, n)$$

and successively deduce $\sum_h d_h c_{hj} = 0$ for $j = n-1, \dots, 1$. That would mean $\sum_h d_h \mathbf{w}_h = 0$, but the $\mathbf{w}_h$ are linearly independent.

By replacing the $\mathbf{w}_h$ with suitable $\mathcal{O}_0$-linear combinations, we can ensure that the c_{hn} are in R and their reductions modulo π are linearly independent over $\mathbb{F}_q$. Now on one hand, the reduction of $(\lambda^i / a_{n-i}^{\sigma^{i-1}}) f(x)$ modulo π is a polynomial in x of the form $b_{i+m} x^{q^{i+m}} + \cdots + b_i x^{q^i}$, which has only q^m distinct roots in $R/\pi R$. On the other hand, the $\mathbb{F}_q$-linear combinations of the reductions of the c_{hn} yield q^{m+1} distinct roots in $R/\pi R$, contradiction.

We conclude that the multiplicity of s as a slope of M is at most m; this establishes the first assertion. To establish the second, note that if $m > 0$, then there exists $c_n \neq 0$ such that $f(c_n) = 0$ by Proposition 3.17; letting c_n be this root, one can then solve for $c_{n-1}, \dots, c_1$ and produce an eigenvector $\mathbf{v}$ with $F\mathbf{v} = \lambda \mathbf{v}$. □

Using this lemma, we can establish the Dieudonné-Manin classification theorem (for which see also Katz [Ka]). We first state it not quite in the standard form. Note: a "basis up to isogeny" means a maximal linearly independent set.

Proposition 5.5. *Suppose R is complete with algebraically closed residue field. Then every σ-module M over R has a basis up to isogeny of eigenvectors of nonnegative slopes over $R \otimes_{\mathcal{O}} \mathcal{O}'$ for some finite extension $\mathcal{O}'$ of $\mathcal{O}$ (depending on M).*

Proof. We proceed by induction on $n = \operatorname{rank} M$. Let $\mathbf{v}$ be any nonzero element of M, and let m be the smallest integer such that $\mathbf{v}, F\mathbf{v}, \dots, F^m\mathbf{v}$ are linearly dependent. Then $N = \operatorname{SatSpan}(\mathbf{v}, F\mathbf{v}, \dots, F^{m-1}\mathbf{v})$ is a σ-submodule of M, and Lemma 5.4 implies that it has a primitive eigenvector $\mathbf{v}_1$ of nonnegative slope over $R \otimes_{\mathcal{O}} \mathcal{O}'$ for some $\mathcal{O}'$ (since the corresponding polynomial has a root of nonnegative valuation there). By the induction hypothesis, we can choose $\mathbf{w}_2, \dots, \mathbf{w}_n$ over $R \otimes_{\mathcal{O}} \mathcal{O}''$ for some $\mathcal{O}''$, whose images in $M/\operatorname{SatSpan}(\mathbf{v}_1)$ form a basis up to isogeny of eigenvectors of nonnegative slopes. We then have $F\mathbf{v}_1 = \lambda_1 \mathbf{v}_1$, where we may take λ_1 fixed by σ, and $F\mathbf{w}_i = \lambda_i \mathbf{w}_i + c_i \mathbf{v}_1$ for some $\lambda_i \in \mathcal{O}$ and $c_i \in R$. Apply Proposition 3.17 to find $a_i \in R$ such that $\lambda_1 c_i + \lambda_1 a_i^{\sigma} - \lambda_i a_i = 0$, and set $\mathbf{v}_i = \lambda_1 \mathbf{w}_i + a_i \mathbf{v}_1$ for $i = 2, \dots, n$; then $F\mathbf{v}_i = \lambda_i \mathbf{v}_i$, so $\mathbf{v}_1, \dots, \mathbf{v}_n$ form a basis up to isogeny of eigenvectors of nonnegative slope over $R \otimes_{\mathcal{O}} \mathcal{O}''$, as desired. □

From this statement we deduce the Dieudonné-Manin classification theorem in its more standard form.

Theorem 5.6 (Dieudonné-Manin). *Suppose R is complete with algebraically closed residue field. Then every σ-module over R is canonically isogenous to the direct sum of σ-modules, each with a single slope, with all of these slopes distinct. Moreover, every σ-module of a single slope is isogenous to a direct sum of standard σ-modules of that slope.*

Proof. Let M be a σ-module over R. For each slope s that occurs in a basis up to isogeny of eigenvectors produced by Proposition 5.5 over $R \otimes_{\mathcal{O}} \mathcal{O}'$, let M_s be the span of all eigenvectors of M over $R \otimes_{\mathcal{O}} \mathcal{O}'$ of slope s. Then M_s is invariant under $\mathrm{Gal}(\mathcal{O}'/\mathcal{O})$, so by Galois descent, M_s descends to a σ-submodule of M. Moreover, M_s is isogenous to a direct sum of standard σ-modules of slope s by Proposition 5.3. This proves the desired result. $\square$

Given a σ-module M over a discrete valuation ring R unramified over $\mathcal{O}$, we can embed R into a complete discrete valuation ring over which M has a basis up to isogeny of eigenvectors by Proposition 5.5. (First complete R, then take its maximal unramified extension, then complete again, then tensor with a suitable $\mathcal{O}'$ over $\mathcal{O}$.) By Proposition 5.2, the slopes and multiplicities do not depend on the choice of the basis. Define the *generic slopes* of M as the slopes of the eigenvectors in the basis, and the *generic Newton polygon* of M as the polygon with vertices (i, y_i) for $i = 0, \dots, \operatorname{rank} M$, where $y_0 = 0$ and $y_i - y_{i-1}$ is the i-th smallest generic slope of M (counting multiplicity). If M has all slopes equal to 0, we say M is *unit-root*.

With the definition of the generic Newton polygon in hand, we can refine the conclusion of Lemma 5.4 as follows.

Proposition 5.7. *Given elements $a_0, \dots, a_{n-1}$ of R with a_0 nonzero, let M be the σ-module with basis $\mathbf{v}_1, \dots, \mathbf{v}_n$ such that*

$$\begin{aligned} F\mathbf{v}_i &= \mathbf{v}_{i+1} \qquad (i = 1, \dots, n-1) \\ F\mathbf{v}_n &= a_0\mathbf{v}_1 + \cdots + a_{n-1}\mathbf{v}_n. \end{aligned}$$

Then the generic Newton polygon of M coincides with the the Newton polygon of the polynomial $x^n + a_{n-1}x^{n-1} + \cdots + a_0$ over R.

Proof. The two Newton polygons have the same length n, and every number occurs at least as often as a slope of the polynomial as it occurs as a slope of M by Lemma 5.4. Thus all multiplicities must coincide. $\square$

For our purposes, the principal consequence of this fact is the following.

Proposition 5.8. *Let M be a σ-module over $R[\frac{1}{p}]$ with all slopes nonnegative. Then M is isomorphic to a σ-module defined over R.*

Proof. We proceed by induction on $n = \operatorname{rank} M$. Let $\mathbf{v} \in M$ be nonzero, and let m be the smallest integer such that $\mathbf{v}, F\mathbf{v}, \dots, F^m\mathbf{v}$ are linearly dependent. Then $F^m\mathbf{v} = a_0\mathbf{v} + \cdots + a_{m-1}F^{m-1}\mathbf{v}$ for some $a_0, \dots, a_{m-1} \in R[\frac{1}{p}]$; by Proposition 5.7, the a_i belong to R. Let $N = \operatorname{SatSpan}(\mathbf{v}, F\mathbf{v}, \dots, F^{m-1}\mathbf{v})$; by the induction hypothesis, M/N is isomorphic to a σ-module defined over R. So we can choose $\mathbf{w}_1, \dots, \mathbf{w}_{n-m}$ that form a basis of M together with $\mathbf{v}, F\mathbf{v}, \dots, F^{m-1}\mathbf{v}$, such that for $i = 1, \dots, n-m$, $F\mathbf{w}_i$ equals an $R[\frac{1}{p}]$-linear combination of the $F^j\mathbf{v}$ plus an R-linear combination of the $\mathbf{w}_j$. For λ sufficiently divisible by π, the basis $\lambda\mathbf{v}, \lambda F\mathbf{v}, \dots, \lambda F^{m-1}\mathbf{v}, \mathbf{w}_1, \dots, \mathbf{w}_{n-m}$ has the property that the image of each basis vector under Frobenius is an R-linear combination of basis vectors. This gives the desired isomorphism. $\square$

We close the section with another method for reading off the generic Newton polygon of a σ-module, inspired by an observation of Buzzard and Calegari [BC, Lemma 5]. (We suspect it may date back earlier, possibly to Manin.)

Proposition 5.9. *Let M be a σ-module over a discrete valuation ring R. Suppose M has a basis on which F acts by the matrix A, where AD^{-1} is congruent to the identity matrix modulo π for some diagonal matrix D over $\mathcal{O}$. Then the slopes of the generic Newton polygon of M equal the valuations of the diagonal entries of D.*

Proof. Without loss of generality we may assume R is complete with algebraically closed residue field. We produce a sequence of matrices $\{U_l\}_{l=1}^\infty$ such that $U_1 = I$, $U_{l+1} \equiv U_l \pmod{\pi^l}$ and $U_l^{-1}AU_l^\sigma D^{-1} \equiv I \pmod{\pi^l}$; the π-adic limit U of the U_l will satisfy $AU^\sigma = UD$, proving the proposition. The conditions for $l = 1$ are satisfied by the assumption that $AD^{-1} \equiv I \pmod{\pi}$.

Suppose U_l has been defined. Put $V = U_l^{-1}AU_l^\sigma D^{-1} - I$. Define a matrix W whose entry W_{ij}, for each i and j, is a solution of the equation $W_{ij} - D_{ii}W_{ij}^\sigma D_{jj}^{-1} = V_{ij}$ with $\min\{v_p(W_{ij}), v_p(D_{ii}W_{ij}^\sigma D_{jj}^{-1})\} = v_p(V_{ij})$ (such a solution exists by Proposition 3.17). Then W and $DW^\sigma D^{-1}$ are both congruent to 0 modulo π^l. Put $U_{l+1} = U_l(I+W)$; then

$$\begin{aligned}
U_{l+1}^{-1}AU_{l+1}^\sigma D^{-1} &= (I+W)^{-1}U_l^{-1}AU_l^\sigma(I+W)^\sigma D^{-1} \\
&= (I+W)^{-1}U_l^{-1}AU_l^\sigma D^{-1}(I + DW^\sigma D^{-1}) \\
&= (I+W)^{-1}(I+V)(I+DW^\sigma D^{-1}) \\
&\equiv I - W + V + DW^\sigma D^{-1} = I \pmod{\pi^{l+1}}.
\end{aligned}$$

Thus the conditions for U_{l+1} are satisfied, and the proposition follows. $\square$

5.3 Slope filtrations

The Dieudonné-Manin classification holds over Γ^K only if K is algebraically closed, and even then does not descend to Γ^K_{con} in general. In this section, we exhibit two partial versions of the classification that hold with weaker conditions on the coefficient ring. One (the descending filtration) is due to de Jong [dJ, Proposition 5.8]; for symmetry, we present independent proofs of both results.

The following filtration result applies for any K but does not descend to Γ_{con}.

Proposition 5.10 (Ascending generic filtration). *Let K be a valued field. Then any σ-module M over $\Gamma = \Gamma^K$ admits a unique filtration $M_0 = 0 \subset M_1 \subset \cdots \subset M_m = M$ by σ-submodules such that*

1. *for $i = 1, \dots, m$, M_{i-1} is saturated in M_i and M_i/M_{i-1} has all generic slopes equal to s_i, and*
2. *$s_1 < \cdots < s_m$.*

Moreover, if K is separably closed and k is algebraically closed, each M_i/M_{i-1} is isogenous to a direct sum of standard σ-modules.

Warning: this proof uses the object Γ^{sep} even though this has only so far been defined for k perfect. Thus we must give an *ad hoc* definition here. For any finite separable extension L over K, Lemma 3.1 produces a finite extension of Γ^K with residue field L, and Lemma 3.2 allows us to identify that extension as a subring of Γ^{alg}. We define Γ^{sep} as the completed union of these subrings; note that $\Gamma^{\mathrm{perf}} \cap \Gamma^{\mathrm{sep}} = \Gamma$.

Proof. By the Dieudonné-Manin classification (Theorem 5.6), M is canonically isogenous to a direct sum of σ-submodules, each of a different single slope. By Corollary 2.7, these submodules descend to Γ^{perf}; let M_1 be the submodule of minimum slope. It suffices to show that M_1 is defined over Γ, as an induction on rank will then yield the general result. Moreover, it is enough to establish this when M_1 has rank 1: if M_1 has rank d, then the lowest slope submodule of $\wedge^d M$ is the rank one submodule $\wedge^d M_1$, and if $\wedge^d M_1$ is defined over Γ, then so is M_1.

So suppose that M_1 has rank 1; this implies that the lowest slope of M belongs to the value group of $\mathcal{O}$. By applying an isogeny and then twisting, we may reduce to the case where the lowest slope is 0. Let $\mathbf{e}_1, \dots, \mathbf{e}_n$ be a basis of M and let A be the matrix such that $F\mathbf{e}_j = \sum_{jl} A_{jl}\mathbf{e}_l$.

Let $\mathbf{v}$ be an eigenvector of M over Γ^{alg} with $F\mathbf{v} = \mathbf{v}$. We will show that $\mathbf{v}$ is congruent to an element of $M \otimes_\Gamma \Gamma^{\mathrm{sep}}$ modulo π^m for each m, by induction on m. The case $m = 0$ is vacuous, so assume the result is known for some m, that is, $\mathbf{v} = \mathbf{w} + \pi^m \mathbf{x}$ with $\mathbf{w} \in M \otimes_\Gamma \Gamma^{\mathrm{sep}}$ and $\mathbf{x} \in M \otimes_\Gamma \Gamma^{\mathrm{alg}}$. Then $0 = F\mathbf{v} - \mathbf{v} = (F\mathbf{w} - \mathbf{w}) + \pi^m(F\mathbf{x} - \mathbf{x})$; that is, $F\mathbf{x} - \mathbf{x}$ belongs to $M \otimes_\Gamma \Gamma^{\mathrm{sep}}$. Write $\mathbf{x} = \sum_j c_j \mathbf{e}_j$ and $F\mathbf{x} - \mathbf{x} = \sum d_j \mathbf{e}_j$, and let s be the smallest nonnegative integer such that the reduction of c_j modulo π lies in K^{1/q^s} for all j. Then $d_j = -c_j + \sum_l A_{jl} c_l^\sigma$; if $s > 0$, then writing $c_j = -d_j + \sum_l A_{jl} c_l^\sigma$ shows that the reduction of c_j lies in $K^{1/q^{s-1}}$ for all j, contradiction. Thus $s = 0$, and $\mathbf{x}$ is congruent modulo π to an element of $M \otimes_\Gamma \Gamma^{\mathrm{sep}}$, completing the induction.

We conclude that $\mathbf{v} \in M \otimes_\Gamma \Gamma^{\mathrm{sep}}$. Thus M_1 is defined both over Γ^{perf} and over Γ^{sep}, so it is in fact defined over $\Gamma^{\mathrm{perf}} \cap \Gamma^{\mathrm{sep}} = \Gamma$, as desired. This proves the desired result, except for the final assertion. In case K is separably closed, one can repeat the above argument over a suitable finite extension of $\mathcal{O}$ to show that each M_i/M_{i-1} is spanned by eigenvectors, then apply Proposition 5.3. $\square$

The following filtration result applies over Γ_{con}, not just over Γ, but requires that K be perfect.

Proposition 5.11 (Descending generic filtration). *Let K be a perfect valued field over k. Then any σ-module M over $\Gamma_{\mathrm{con}} = \Gamma^K_{\mathrm{con}}$ admits a unique filtration $M_0 = 0 \subset M_1 \subset \cdots \subset M_m = M$ by σ-submodules such that*

1. *for $i = 1, \dots, m$, M_{i-1} is saturated in M_i and M_i/M_{i-1} has all generic slopes equal to s_i, and*

2. *$s_1 > \cdots > s_m$.*

Moreover, if K is algebraically closed, each M_i/M_{i-1} is isogenous to a direct sum of standard σ-modules.

Proof. By the Dieudonné-Manin classification (Theorem 5.6), M is canonically isogenous to a direct sum of σ-submodules, each of a different single slope. By Corollary 2.7, these submodules descend to Γ; let M_1 be the submodule of maximum slope. It suffices to show that M_1 is defined over Γ_{con}, as an induction on rank will then yield the general result. Moreover, it is enough to establish this when M_1 has rank 1: if M_1 has rank d, then the lowest slope submodule of $\wedge^d M$ is the rank one submodule $\wedge^d M_1$, and if $\wedge^d M_1$ is defined over Γ_{con}, then so is M_1.

So suppose that M_1 has rank 1; this implies that the highest slope of M belongs to the value group of $\mathcal{O}$. Choose $\lambda \in \mathcal{O}$ whose valuation equals that slope. Let $\mathbf{v}_1, \dots, \mathbf{v}_n$ be a basis of $M \otimes_{\Gamma_{\mathrm{con}}} \Gamma^{\mathrm{alg}}$, in which $F\mathbf{v}_1 = \lambda \mathbf{v}_1$ and the remaining $\mathbf{v}_i$ span the submodules of M of lower slopes. Choose $\mathbf{w}_i \in M \otimes_{\Gamma_{\mathrm{con}}} \Gamma^{\mathrm{alg}}_{\mathrm{con}}$ sufficiently close π-adically to $\mathbf{v}_i$ for $i = 1, \dots, n$ so that the matrix B with $\lambda \mathbf{w}_i = \sum_j B_{ij} F\mathbf{w}_j$ has entries in Γ^{alg} and

$$B_{ij} \equiv \begin{cases} 1 & i = j = 1 \\ 0 & \text{otherwise} \end{cases} \pmod{\pi};$$

this is possible because the congruence holds for $\mathbf{w}_i = \mathbf{v}_i$. Then the $\mathbf{w}_i$ form a basis of $M \otimes_{\Gamma_{\mathrm{con}}} \Gamma^{\mathrm{alg}}_{\mathrm{con}}$.

Write $\mathbf{v}_1 = \sum_i c_i \mathbf{w}_i$, so that $c_i^\sigma = \sum_j B_{ji} c_j$. Since $v_0(B) \geq 0$, we can find r such that $w_r(B) \geq 0$. We now show that $r v_h(c_i) + h \geq 0$ for all i and h, by induction on h. Suppose this holds with h replaced by any smaller value. Then the equality $c_i^\sigma = \sum_j B_{ji} c_j$ implies

$$q v_h(c_i) \geq \min_{l,j} \{ v_l(B_{ji}) + v_{h-l}(c_j) \}.$$

Choose j, l for which the minimum is achieved. If $l = 0$, then we must have $i = j = 1$, in which case $v_0(B_{11}) = 0$ and $q v_h(c_1) \geq v_h(c_1)$, whence $v_h(c_1) \geq 0$ and $r v_h(c_1) + h \geq 0$ as well.

If the minimum occurs for some $l > 0$, then

$$\begin{aligned} rv_h(c_i) + h &\geq rq^{-1}(v_l(B_{ji}) + v_{h-l}(c_j)) + h \\ &\geq rq^{-1}(v_l(B_{ji}) + v_{h-l}(c_j)) + q^{-1}h \\ &\geq q^{-1}(rv_l(B_{ji}) + l + rv_{h-l}(c_j) + (h-l)) \\ &\geq q^{-1}(0+0) = 0 \end{aligned}$$

by the induction hypothesis. Therefore $rv_h(c_i) + h \geq 0$ for all h, so $c_i \in \Gamma^{\mathrm{alg}}_{\mathrm{con}}$ for each i.

We conclude that $\mathbf{v}_1 \in M \otimes_\Gamma \Gamma^{\mathrm{alg}}_{\mathrm{con}}$. Thus M_1 is defined both over Γ and over $\Gamma^{\mathrm{alg}}_{\mathrm{con}}$, so it is in fact defined over $\Gamma \cap \Gamma^{\mathrm{alg}}_{\mathrm{con}} = \Gamma_{\mathrm{con}}$, as desired. This proves the desired result, except for the final assertion. In case K is algebraically closed, one can repeat the above argument over a suitable finite extension of $\mathcal{O}$ to show that each M_i/M_{i-1} is spanned by eigenvectors, then apply Proposition 5.3. □

Although we will not use the following result explicitly, it is worth pointing out.

Corollary 5.12. *Let K be a valued field, for k algebraically closed. Then any σ-module M over Γ^K_{con}, all of whose generic slopes are equal, is isogenous over $\Gamma^{\mathrm{sep}}_{\mathrm{con}}$ to a direct sum of standard σ-modules.*

Proof. In this case, the ascending and descending filtrations coincide, so both are defined over $\Gamma^K \cap \Gamma^{\mathrm{perf}}_{\mathrm{con}} = \Gamma^K_{\mathrm{con}}$ and the eigenvectors are defined over $\Gamma^{\mathrm{sep}}_{\mathrm{con}} \otimes_{\mathcal{O}} \mathcal{O}'$ for some finite extension $\mathcal{O}'$ of $\mathcal{O}$. Thus the claim follows from Proposition 5.3. □

5.4 Comparison of the Newton polygons

A σ-module over Γ_{con} can be base-extended both to Γ and to $\Gamma_{\mathrm{an,con}}$; as a result, it admits both a generic and a special Newton polygon. In this section, we compare these two polygons. The main results are that the special polygon lies above the generic polygon, and that when the two coincide, the σ-module admits a partial decomposition over Γ_{con} (reminiscent of the Newton-Hodge decomposition of [Ka]).

Throughout this section, K is an arbitrary valued field, which we suppress from the notation.

Proposition 5.13. *Let M and N be σ-modules over Γ_{con}. Let $r_1, \dots, r_m$ and $s_1, \dots, s_n$ be the generic (resp. special) slopes of M and N.*

1. *The generic (resp. special) slopes of $M \oplus N$ are $r_1, \dots, r_m, s_1, \dots, s_n$.*
2. *The generic (resp. special) slopes of $M \otimes N$ are $r_i + s_j$ for $i = 1, \dots, m$ and $j = 1, \dots, n$.*
3. *The generic (resp. special) slopes of $\wedge^l M$ are $r_{i_1} + \cdots + r_{i_l}$ for $1 \leq i_1 < \cdots < i_l \leq m$.*
4. *The generic (resp. special) slopes of M^* are $-r_1, \dots, -r_m$.*

Proof. These results follow immediately from the definition of the generic (resp. special) Newton slopes as the valuations of the eigenvalues of a basis of eigenvectors of M over Γ^{alg} (resp. $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$). $\square$

Proposition 5.14. *Let M be a σ-module over Γ_{con}. Then the special Newton polygon lies above the generic Newton polygon, and both have the same endpoint.*

Proof. The Newton polygons coincide for M of rank 1 because M has an eigenvector over $\Gamma^{\mathrm{alg}}_{\mathrm{con}}$ by Proposition 3.18. Thus the Newton polygons of $\wedge^n M$ coincide for $n = \operatorname{rank} M$; that is, the Newton polygons of M have the same endpoint. By the descending slope filtration (Proposition 5.11), M admits a basis $\mathbf{w}_1, \dots, \mathbf{w}_n$ over $\Gamma^{\mathrm{alg}}_{\mathrm{con}}$ such that modulo $\mathrm{SatSpan}(\mathbf{w}_1, \dots, \mathbf{w}_{i-1})$, $\mathbf{w}_i$ is an eigenvector whose slope is the i-th largest generic slope of M. Let $\mathbf{v}_1, \dots, \mathbf{v}_n$ be a basis of eigenvectors of M over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$; then by Proposition 4.5, the sequence of valuations of the eigenvalues of the $\mathbf{w}_i$ majorizes that of the $\mathbf{v}_i$. In other words, the sequence of generic slopes majorizes the sequence of special slopes, whence the comparison of Newton polygons. $\square$

Proposition 5.15. *Let $0 \to M_1 \to M \to M_2 \to 0$ be an exact sequence of σ-modules over Γ_{con}. Suppose the least generic slope of M_2 is greater than the greatest generic slope of M_1. Then the special Newton polygon of M is equal to the union of the special Newton polygons of M_1 and M_2.*

Proof. The least generic slope of M_2 is less than or equal to its least special slope, and the greatest generic slope of M_1 is greater than or equal to its greatest special slope, both by Proposition 5.14. Thus we may apply Proposition 4.3 over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ (after extending $\mathcal{O}$ suitably) to deduce the desired result. $\square$

It is perhaps not surprising that when the generic and special Newton polygons coincide, one gets a slope filtration that descends further than usual.

Proposition 5.16. *Let M be a σ-module over Γ_{con} whose generic and special Newton polygons coincide. Then M admits an ascending slope filtration over Γ_{con}.*

Proof. We need to show that the ascending slope filtration of Proposition 5.10 is defined over Γ_{con}; it is enough to verify this after enlarging $\mathcal{O}$. This lets us assume that k is algebraically closed, and that the value group of $\mathcal{O}$ contains all of the slopes of M. By Theorem 4.16, we can find a basis $\mathbf{v}_1, \dots, \mathbf{v}_n$ of eigenvectors of M over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$, with $F\mathbf{v}_i = \lambda_i \mathbf{v}_i$ for $\lambda_i \in \mathcal{O}_0[\frac{1}{p}]$ such that $v_p(\lambda_1) \geq \cdots \geq v_p(\lambda_n)$. By Proposition 5.11 (the descending slope filtration), we can find a basis up to isogeny $\mathbf{w}_1, \dots, \mathbf{w}_n$ of M over $\Gamma^{\mathrm{alg}}_{\mathrm{con}}$ such that $F\mathbf{w}_i = \lambda_i \mathbf{w}_i + \sum_{j<i} A_{ij} \mathbf{w}_j$ for some $A_{ij} \in \Gamma^{\mathrm{alg}}_{\mathrm{con}}$.

Write $\mathbf{v}_n = \sum_i b_i \mathbf{w}_i$ with $b_i \in \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$; applying F to both sides, we have $\lambda_n b_i = \lambda_i b_i^\sigma + \sum_{j>i} b_j^\sigma A_{ji}$ for $i = 1, \dots, n$. By Proposition 3.19(a) and (c), we obtain $b_i \in \Gamma^{\mathrm{alg}}_{\mathrm{con}}[\frac{1}{p}]$ for $i = n, n-1, \dots, 1$, and so $\mathbf{v}_n$ is defined over $\Gamma^{\mathrm{alg}}_{\mathrm{con}}[\frac{1}{p}]$.

By repeating the above reasoning, we see that the image of $\mathbf{v}_i$ in $M/\mathrm{SatSpan}(\mathbf{v}_{i+1}, \dots, \mathbf{v}_n)$ is defined over $\Gamma^{\mathrm{alg}}_{\mathrm{con}}[\frac{1}{p}]$ for $i = n, \dots, 1$. Thus the ascending slope filtration is defined over $\Gamma^{\mathrm{alg}}_{\mathrm{con}}[\frac{1}{p}]$. Since it is also defined over Γ by Proposition 5.10, it is in fact defined over $\Gamma \cap \Gamma^{\mathrm{alg}}_{\mathrm{con}}[\frac{1}{p}] = \Gamma_{\mathrm{con}}$, as desired. $\square$

6 From a slope filtration to quasi-unipotence

In this chapter we construct a canonical filtration of a σ-module over $\Gamma^{k((t))}_{\mathrm{an,con}}$. We do this by partially descending the special slope filtration obtained over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ in Chapter 4. More specifically, we show that by changing basis over a finite extension of $\Gamma_{\mathrm{an,con}}$, we can make Frobenius act by a matrix with entries in a finite extension of Γ_{con}, whose generic Newton polygon coincides with the special Newton polygon, allowing the use of Proposition 5.16. This will yield the desired filtration (Theorem 6.10), from which we deduce the p-adic local monodromy theorem (Theorem 1.1) using the quasi-unipotence of unit-root (σ, ∇)-modules over Γ_{con}; the latter is a theorem of Tsuzuki [T1] (for which see also Christol [Ch]).

6.1 Approximation of matrices

We collect some results that allow us to approximate matrices from a large ring with matrices from smaller rings. Note: we will need the notions of slopes and Newton polygons from Section 3.5.

Lemma 6.1. *Let K be a finite extension of $k((t))$ and suppose Γ^K_r contains a unit lifting a uniformizer of K. Then for any $x, y \in \Gamma^K_r[\frac{1}{p}]$, x is coprime to $y + \pi^j$ for all sufficiently large integers j.*

Proof. Suppose on the contrary that x and $y + \pi^j$ fail to be coprime for $j = j_1, j_2, \dots$. By Corollary 3.32, the ideal $(x, y + \pi^{j_l})$ in $\Gamma^K_r[\frac{1}{p}]$ is principal; let d_l be a generator. Note that $(y + \pi^{j_i}, y + \pi^{j_l})$ contains the unit $\pi^{j_i} - \pi^{j_l}$ for $i \neq l$, so is the unit ideal; that means the d_l are pairwise coprime, and x is divisible by $d_1 \cdots d_l$ for any l. But x has only finite total multiplicity while each d_l has nonzero total multiplicity, contradiction. Hence x is coprime to $y + \pi^j$ for j sufficiently large, as desired. $\square$

By an *elementary operation* on a matrix over a ring, we mean one of the following operations:

(a) adding a multiple of one row to another;

(b) multiplying one row by a unit of the ring;

(c) interchanging two rows.

An *elementary matrix* is one obtained from the identity matrix by a single elementary operation; multiplying a matrix on the right by an elementary matrix has the same effect as performing the corresponding elementary operation.

Lemma 6.2. *Pick s such that $0 < s < r$, and let U be a matrix over $\Gamma^{\mathrm{alg}}_{\mathrm{an},r}$ such that $w_l(\det(U) - 1) > 0$ for $s \leq l \leq r$. Then there exists an invertible matrix V over $\Gamma^K_r[\frac{1}{p}]$, for some finite extension K of $k((t))$, such that $w_l(UV - I) > 0$ for $s \leq l \leq r$. Moreover, if U is defined over $\Gamma^{k((t))}_{\mathrm{an},r}$ and t lifts to a semiunit in $\Gamma^{k((t))}_r$, then we may take $K = k((t))$.*

Although we only will apply this when U is invertible, we need to formulate the more general statement in order to carry out the induction.

Proof. We induct on n, the case $n = 1$ being vacuous. Let M_i denote the cofactor of U_{ni} in U, so that $\det(U) = \sum_i M_i U_{ni}$; note that $M_i = (U^{-1})_{in} \det(U)$ in $\mathrm{Frac}(\Gamma^{\mathrm{alg}}_{\mathrm{an},r})$. Let d be a generator of the ideal $(M_1, \dots, M_n)$ in $\Gamma^{\mathrm{alg}}_{\mathrm{an},r}$. Then d divides $\det(U)$; by the hypothesis that $w_l(\det(U) - 1) > 0$ for $s \leq l \leq r$, the largest slope of $\det(U)$ is less than s, so the largest slope of d is also less than s. By Lemma 3.24, there exists a unit $u \in \Gamma^{\mathrm{alg}}_{\mathrm{an},r}$ such that $w_l(ud - 1) > 0$ for $s \leq l \leq r$.

Let $\alpha_1, \dots, \alpha_n$ be elements of $\Gamma^{\mathrm{alg}}_{\mathrm{an},r}$ such that $\sum_i \alpha_i M_i = ud$. Choose $\beta_1, \dots, \beta_{n-1}$ and $\beta'_n \in \Gamma^L_r[\frac{1}{p}]$, for some finite extension L of $k((t))$, so that for $s \leq l \leq r$,

$$w_l(\beta_i - \alpha_i) > 0 \quad (i = 1, \dots, n-1), \qquad \text{and} \qquad w_l(\beta'_n - \alpha_n) > 0.$$

By Lemma 6.1, we can find j for which $\beta_n = \beta'_n + \pi^j$ has the properties that $w_l(\beta_n - \alpha_n) > 0$ for $s \leq l \leq r$ and $(\beta_1, \dots, \beta_n)$ is the unit ideal in $\Gamma^L_r[\frac{1}{p}]$. (Both hold for j sufficiently large.)

By Corollary 3.32, $\Gamma^L_r[\frac{1}{p}]$ is a Bézout ring. Thus Lemma 2.3 can be applied to produce a matrix A over $\Gamma^L_r[\frac{1}{p}]$ of determinant 1 such that $A_{ni} = \beta_i$ for $i = 1, \dots, n$. Put $U' = UA^{-1}$, and let M'_n be the cofactor of U'_{nn} in U'. Then

$$\begin{aligned} M'_n &= ((U')^{-1})_{nn} \det(U') \\ &= (AU^{-1})_{nn} \det(U) \det(A^{-1}) \\ &= \sum_i A_{ni} (U^{-1})_{in} \det(U) \\ &= \sum_i \beta_i M_i, \end{aligned}$$

so that

$$M'_n - 1 = ud - 1 + \sum_i (\beta_i - \alpha_i) M_i$$

and hence $w_l(M'_n - 1) > 0$ for $s \leq l \leq r$.

Apply the induction hypothesis to the upper left $(n-1) \times (n-1)$ submatrix of U', let V' be the resulting matrix, and enlarge L if needed so that V' has entries in $\Gamma^L_{\mathrm{an},r}$. Extend V' to an $n \times n$ matrix by setting $V'_{nn} = 1$ and $V'_{ni} = V'_{in} = 0$ for $i = 1, \dots, n-1$. Then for $s \leq l \leq r$, $w_l((U'V' - I)_{ij}) > 0$ for $1 \leq i, j \leq n-1$. Moreover, $w_l(\det(V') - 1) > 0$, so $w_l(\det(U'V') - 1) > 0$ as well.

We now exhibit a sequence of elementary operations which can be performed on $U'V'$ to obtain a new matrix W over $\Gamma^{\mathrm{alg}}_{\mathrm{an},r}$ with $w_l(W - I) > 0$ for $s \leq l \leq r$; it may clarify matters to regard the procedure as an “approximate Gaussian elimination”. First, define a sequence of matrices $\{X^{(h)}\}_{h=0}^{\infty}$ by $X^{(0)} = U'V'$ and

$$X_{ij}^{(h+1)} = \begin{cases} X_{ij}^{(h)} & i < n \\ X_{nj}^{(h)} - \sum_{m=1}^{n-1} X_{nm}^{(h)} X_{mj}^{(h)} & i = n; \end{cases}$$

note that $X^{(h+1)}$ is obtained from $X^{(h)}$ by subtracting $X_{nm}^{(h)}$ times the m-th row from the n-th row for $m = 1, \dots, n-1$. At each step, $\min_{1 \leq j \leq n-1}\{w_l(X_{nj}^{(h)})\}$ increases by at least $\min_{1 \leq i,j \leq n-1}\{w_l((U'V' - I)_{ij})\}$; thus for h sufficiently large, we have

$$w_l(X_{nj}^{(h)}) > \max\left\{0, \max_{1 \leq i \leq n-1}\{-w_l(X_{in}^{(h)})\}\right\} \qquad (s \leq l \leq r;\ j = 1, \dots, n-1).$$

Pick such an h and set $X = X_h$. Then $w_l((X - I)_{ij}) > 0$ for $1 \leq i \leq n$ and $1 \leq j \leq n-1$, $w_l(X_{in}X_{nj}) > 0$ for $1 \leq i, j \leq n-1$, and $w_l(\det(X) - 1) > 0$. These together imply $w_l(X_{nn} - 1) > 0$.

Next, define a sequence of matrices $\{W^{(h)}\}_{h=0}^{\infty}$ by $W^{(0)} = X$ and

$$W_{ij}^{(h+1)} = \begin{cases} W_{ij}^{(h)} - W_{in}^{(h)} W_{nj}^{(h)} & i < n \\ W_{ij}^{(h)} & i = n; \end{cases}$$

note that $W^{(h+1)}$ is obtained from $W^{(h)}$ by subtracting $W_{in}^{(h)}$ times the n-th row from the i-th row for $i = 1, \dots, n-1$. At each step, $w_l(X_{in}^{(h)})$ increases by at least $w_l(X_{nn}^{(h)} - 1)$; thus for h sufficiently large, we have

$$w_l(W_{in}^{(h)}) > 0 \qquad (s \leq l \leq r;\ i = 1, \dots, n-1).$$

Pick such an h and set $W = W_h$; then $w_l(W - I) > 0$ for $s \leq l \leq r$.

To conclude, note that by construction, $(U'V')^{-1}W$ is a product of elementary matrices over $\Gamma_{\mathrm{an},r}^{\mathrm{alg}}$ of type (a). By suitably approximating each elementary matrix by one defined over $\Gamma_r^K[\frac{1}{p}]$ for a suitable finite extension K of L, we get a matrix X such that $w_l(U'V'X - I) > 0$ for $s \leq l \leq r$. We may thus take $V = A^{-1}V'X$. □

We will need a refinement of the above result.

Lemma 6.3. *Pick s such that $0 < s < r$, and let U be a matrix over $\Gamma_{\mathrm{an},r}^{\mathrm{alg}}$ such that $w_l(\det(U) - 1) > 0$ for $s \leq l \leq r$. Then for any $c > 0$, there exists a finite extension K of $k((t))$ and an invertible matrix V over $\Gamma_r^K[\frac{1}{p}]$ such that $w_l(UV - I) \geq c$ for $s \leq l \leq r$. Moreover, if U is defined over $\Gamma_{\mathrm{an},r}^{k((t))}$ and t lifts to a semiunit in $\Gamma_r^{k((t))}$, then we may take $K = k((t))$.*

Proof. Put

$$s' = s(1 + c/v_p(\pi))^{-1}.$$

Apply Lemma 6.2 to obtain a finite extension L of $k((t))$ and an invertible matrix V' over $\Gamma_r^L[\frac{1}{p}]$ (with $L = k((t))$ in case U is defined over $\Gamma_{\mathrm{an},r}^{k((t))}$) such that $w_l(UV' - I) > 0$ for $s' \leq l \leq r$.

Choose semiunit decompositions $\sum_h W_{ijh}\pi^h$ of $(UV')_{ij} - I$ for $1 \leq i, j \leq n$. For $s \leq l \leq r$ and $m < 0$ in the value group of $\mathcal{O}$, we deduce from $w_{s'}(UV' - I) > 0$ that

$$\begin{aligned} lv_m(UV' - I) + m &= (l/s')(s'v_m(UV' - I) + m) - m(l/s' - 1) \\ &> -m(l/s' - 1) \\ &> v_p(\pi)(s/s' - 1) \\ &= c. \end{aligned}$$

Define a matrix X by $X_{ij} = \sum_{h \geq 0} W_{ijh}\pi^h$; then $UV' - I - X = \sum_{h<0} W_{ijh}\pi^h$, so that for $s \leq l \leq r$,

$$w_l(UV' - I - X) = \min_{m<0}\{lv_m(UV' - I) + m\} \geq c.$$

By construction, $v_m(X) = \infty$ for $m < 0$ and $v_0(X) > 0$. Thus $I + X$ is invertible over $\Gamma_{\mathrm{an},r}^{\mathrm{alg}}$. Choose a matrix W over Γ_r^K, for some finite extension K of L (with $K = k((t))$ if U is defined over $\Gamma_{\mathrm{an},r}^{k((t))}$), such that $w_l(W - (I + X)^{-1}) \geq c$ for $s \leq l \leq r$. Then W is invertible over Γ_r^K, and for $s \leq l \leq r$,

$$\begin{aligned} w_l(UV'W - I) &= w_l((UV' - I - X)W + (I + X)(W - (I + X)^{-1})) \\ &\geq \min\{w_l(UV' - I - X) + w_l(W), w_l(I + X) + w_l(W - (I + X)^{-1})\} \\ &\geq c. \end{aligned}$$

We may thus take $V = V'W$. □

6.2 Some matrix factorizations

Throughout this section, we take $K = k((t))$ and omit it from the notation; note also the use of the naïve partial valuations. Let Γ_u and $\Gamma_{\mathrm{an},u}$ denote the subrings of Γ_{con} and $\Gamma_{\mathrm{an,con}}$, respectively, consisting of elements x of the form $\sum_{i=0}^{\infty} x_i u^i$.

Lemma 6.4. *For $r > 0$ and $c > 0$, let A be a matrix over $\Gamma_{\mathrm{an},r}$ such that $w_r^{\mathrm{naive}}(A - I) \geq c$. Then there exists a unique pair of matrices $U = I + \sum_{i=1}^{\infty} U_i u^i$ over $\Gamma_{\mathrm{an},r}$ and $V = \sum_{i=0}^{\infty} V_i u^{-i}$ over Γ_r such that $w_r^{\mathrm{naive}}(U - I) > 0$, $w_r^{\mathrm{naive}}(V - I) > 0$, and $A = UV$. Moreover, these matrices satisfy $w_r^{\mathrm{naive}}(U - I) \geq c$ and $w_r^{\mathrm{naive}}(V - I) \geq c$.*

Proof. Define a sequence of matrices $\{B^{(j)}\}_{j=0}^{\infty}$ as follows. Begin by setting $B^{(0)} = I$. Given $B^{(j)}$ for some j, put $A(B^{(j)})^{-1} = \sum_{i=-\infty}^{\infty} X_i^{(j)} u^i$, $C^{(j)} = \sum_{i \leq 0} X_i^{(j)} u^i$, $D^{(j)} = \sum_{i>0} X_i^{(j)} u^i$, and put $B^{(j+1)} = C^{(j)} B^{(j)}$.

Since $w_r^{\mathrm{naive}}(A - I) \geq c$, we have $w_r^{\mathrm{naive}}(C^{(0)} - I) \geq c$ and $w_r^{\mathrm{naive}}(D^{(0)}) \geq c$ as well. Thus $w_r^{\mathrm{naive}}(A(B^{(1)})^{-1} - I) \geq c$, and by induction one has $w_r^{\mathrm{naive}}(C^{(j)} - I) \geq c$ and $w_r^{\mathrm{naive}}(D^{(j)}) \geq c$

for all j. But we can do better, by showing by induction that $w_r^{\text{naive}}(C^{(j)} - I) \geq (j+1)c$ and $w_r^{\text{naive}}(D^{(j+1)} - D^{(j)}) \geq (j+2)c$ for $j \geq 0$. Given $w_r^{\text{naive}}(C^{(j)} - I) \geq (j+1)c$, we have

$$\begin{aligned}A(B^{(j+1)})^{-1} - I &= A(B^{(j)})^{-1}(C^{(j)})^{-1} - I \\ &= (C^{(j)} + D^{(j)})(C^{(j)})^{-1} - I \\ &= D^{(j)}(C^{(j)})^{-1} \\ &= D^{(j)} + D^{(j)}((C^{(j)})^{-1} - I).\end{aligned}$$

Since $D^{(j)}$ has only positive powers of u, $C^{(j+1)}$ is equal to the sum of the terms of $I + D^{(j)}((C^{(j)})^{-1} - I)$ involving nonpositive powers of u. In particular,

$$w_r^{\text{naive}}(C^{(j+1)} - I) \geq w_r^{\text{naive}}(D^{(j)}((C^{(j)})^{-1} - I)) \geq c + (j+1)c = (j+2)c;$$

likewise, $D^{(j+1)} - D^{(j)}$ consists of terms from $D^{(j)}((C^{(j)})^{-1} - I)$, so $w_r^{\text{naive}}(D^{(j+1)} - D^{(j)}) \geq (j+2)c$. This completes the induction.

Since $C^{(j)}$ converges to I, we see that $B^{(j)}$ converges to a limit V such that $w_r^{\text{naive}}(V - I) \geq c$. Under w_r^{naive}, $I + D^{(j)}$ also converges to a limit U such that $w_r^{\text{naive}}(U - I) \geq c$, and $A(B^{(j)})^{-1} - I - D^{(j)}$ converges to 0. Therefore $AV^{-1} = U$ has entries in $\Gamma_{\text{an},r,\text{naive}}$, and U and V satisfy the desired conditions.

This establishes the existence of the desired factorization. To establish uniqueness, suppose we have a second decomposition $A = U'V'$ with $U' - I$ only involving positive powers of u, V' only involving negative powers of u, $w_r^{\text{naive}}(U' - I) > 0$, and $w_r^{\text{naive}}(V' - I) > 0$. Within the completion of $\Gamma_r[\frac{1}{p}]$ with respect to $|\cdot|_r$, the matrices U, V, U', V' are invertible and $(U')^{-1}U = V'V^{-1}$. On the other hand, $(U')^{-1}U - I$ involves only positive powers of u, while $V'V^{-1} - I$ involves no positive powers of u. This is only possible if $(U')^{-1}U - I = V'V^{-1} - I = 0$, which yields $U = U'$ and $V = V'$. □

The following proposition may be of interest outside of its use to prove the results of this paper. For example, Berger's proof [Bg, Corollaire 0.3] that any crystalline representation is of finite height uses a lemma from [Ke1] equivalent to this.

Proposition 6.5. *Let $A = \sum_{i=-\infty}^{\infty} A_i u^i$ be an invertible matrix over $\Gamma_{\text{an,con}}$. Then there exist invertible matrices U over $\Gamma_{\text{an},u}$ and V over $\Gamma_{\text{con}}[\frac{1}{p}]$ such that $A = UV$. Moreover, if $w_r^{\text{naive}}(A - I) > 0$ for some $r > 0$, there is a unique choice of U and V such that $U - I$ involves only positive powers of u, V involves no positive powers of u, $w_r^{\text{naive}}(U - I) > 0$ and $w_r^{\text{naive}}(V - I) > 0$; for these U and V, $\min\{w_r^{\text{naive}}(U - I), w_r^{\text{naive}}(V - I)\} \geq w_r^{\text{naive}}(A - I)$.*

Proof. By Lemma 6.2, there exists an invertible matrix W over $\Gamma_{\text{con}}[\frac{1}{p}]$ such that $w_r^{\text{naive}}(AW - I) > 0$. Apply Lemma 6.4 to write $AW = U_1V_1$ for matrices U_1 over $\Gamma_{\text{an},u}$ and V_1 over Γ_{con}, and to write $(AW)^{-T} = U_2V_2$ for matrices U_2 over $\Gamma_{\text{an},u}$ and V_2 over Γ_{con}. Now $I = (AW)^T(AW)^{-T} = V_1^T U_1^T U_2 V_2$, and so $V_1^{-T}V_2^{-1} = U_1^T U_2$ has entries in $\Gamma_{\text{con}} \cap \Gamma_{\text{an},u} = \Gamma_u$. Moreover, $U_1^T U_2 - I$ involves only positive powers of u, so $U_1^T U_2$ is invertible over Γ_u and U_1 is invertible over $\Gamma_{\text{an},u}$. Our desired factorization is now $A = UV$ with $U = U_1$ and $V = V_1W^{-1}$. If $w_r^{\text{naive}}(A - I) > 0$, we may take $W = I$ above and deduce the uniqueness from Lemma 6.4. □

So far we have exhibited factorizations that separate positive and negative powers of u. We use these to give a factorization that separates a matrix over $\Gamma_{\mathrm{an,con}}$ into a matrix over Γ_{con} times a matrix with only positive powers of u, in such a way that the closer the original matrix is to being defined over Γ_{con}, the smaller the positive matrix will be.

Proposition 6.6. *Let A be an invertible matrix over $\Gamma_{\mathrm{an},r}$ such that $w_r^{\mathrm{naive}}(A-I) > 0$. Then there exists a canonical pair of invertible matrices U over $\Gamma_{\mathrm{an},u}$ and V over Γ_{con} such that $A = UV$, $U-I$ has only positive powers of u, $V-I \equiv 0 \pmod{\pi}$, $w_r^{\mathrm{naive}}(V-I) \geq w_r^{\mathrm{naive}}(A-I)$ and*

$$w_r^{\mathrm{naive}}(U-I) \geq \min_{m \leq 0}\{r v_m^{\mathrm{naive}}(A-I) + m\}.$$

Here "canonical" does not mean "unique". It means that the construction of U and V depends only on A and not on r.

Proof. Write $A - I = \sum_i A_i u^i$, and let X be the sum of A_i over all i for which $v_p(A_i) > 0$. Then

$$\begin{aligned} w_r^{\mathrm{naive}}(A(I+X)^{-1} - I) &\geq w_r^{\mathrm{naive}}(A - I - X) + w_r^{\mathrm{naive}}((I+X)^{-1}) \\ &= \min_{v_p(A_i) \leq 0}\{v_p(A_i) + ri\} \\ &= \min_{m \leq 0}\{r v_m^{\mathrm{naive}}(A-I) + m\}. \end{aligned}$$

Apply Proposition 6.5 to factor $A(I+X)^{-1}$ as BC, where

$$\min\{w_r^{\mathrm{naive}}(B-I), w_r^{\mathrm{naive}}(C-I)\} \geq \min_{m \leq 0}\{r v_m^{\mathrm{naive}}(A-I) + m\},$$

$B - I$ involves only positive powers of u, and C involves no positive powers of u; the desired matrices are $U = B$ and $V = C(I+X)$. $\square$

6.3 Descending the special slope filtration

In this section, we refine the decomposition given by Theorem 4.16 in the case of a σ-module defined over $\Gamma_{\mathrm{an,con}}^{k((t))}$, to obtain our main filtration theorem.

Lemma 6.7. *For K a valued field and $r > 0$ satisfying the conclusion of Proposition 3.11, let U be a matrix over $\Gamma_{\mathrm{an},r}^K$ and V a matrix over Γ_r^K such that $w_r(V-I) > 0$ and $v_p(V-I) > 0$. Then*

$$\min_{m \leq 0}\{r v_m(UV - I) + m\} = \min_{m \leq 0}\{r v_m(U-I) + m\}.$$

Proof. In one direction, we have

$$\begin{aligned} \min_{m \leq 0}\{r v_m(UV - I) + m\} &= \min_{m \leq 0}\{r v_m((U-I)V + (V-I)) + m\} \\ &= \min_{m \leq 0}\{r v_m((U-I)V) + m\} \\ &\geq \min_{m \leq 0, l \geq 0}\{r v_l(V) + l + r v_{m-l}(U-I) + (m-l)\} \\ &\geq \min_{m \leq 0}\{r v_m(U-I) + m\}, \end{aligned}$$

the last inequality holding because $w_r(V) = 0$. The reverse direction is implied by the above inequality with U and V replaced by UV and V^{-1}. $\square$

The key calculation is the following proposition. In fact, it should be possible to give a condition of this form that guarantees that a σ-module has a particular special Newton polygon. However, we have not found such a condition so far.

Proposition 6.8. *Let K be a finite extension of $k((t))$ and $r > 0$ a number for which there exists a semiunit u in Γ_{qr}^K lifting a uniformizer of K. Let A be an invertible matrix over $\Gamma_{\mathrm{an},r}$, and suppose that there exists a diagonal matrix D over $\mathcal{O}$ such that*

$$w_r(AD^{-1} - I) > \max_{i,j}\{v_p(D_{ii}) - v_p(D_{jj})\}.$$

Then there exists an invertible matrix U over $\Gamma_{\mathrm{an},qr}$ such that $w_r(U - I) > 0$, $U - I$ involves only positive powers of u, $U^{-1}AU^\sigma D^{-1}$ is invertible over Γ_r and $v_p(U^{-1}AU^\sigma D^{-1} - I) > 0$.

Proof. There is no loss of generality in assuming $K = k((t))$. Then by Lemma 3.7, for $s \leq qr$ and $x \in \Gamma_{\mathrm{an},r}$, $w_s(x) = w_s^{\mathrm{naive}}(x)$ and $\min_{m\leq 0}\{sv_m(x) + m\} = \min_{m\leq 0}\{sv_m^{\mathrm{naive}}(x) + m\}$. This allows us to apply the results of the previous section.

Put $c = \max_{i,j}\{v_p(D_{ii}) - v_p(D_{jj})\}$ and $d = w_r(AD^{-1} - I)$, and define sequences $\{A_i\}$, $\{U_i\}$, $\{V_i\}$ for $i = 0, 1, \dots$ as follows. Begin with $A_0 = A$. Given A_i, factor A_iD^{-1} as U_iV_i as per Proposition 6.6, and set $A_{i+1} = U_i^{-1}A_iU_i^\sigma$, so that $A_{i+1}D^{-1} = V_i(DU_i^\sigma D^{-1})$.

Note that the application of Proposition 6.6 is only valid if $w_r(A_iD^{-1} - I) > 0$. In fact, we will show that

$$\min_{m\leq 0}\{rv_m(A_iD^{-1} - I) + m\} \geq d + i((q-1)d - c) \qquad \text{and} \qquad w_r(A_iD^{-1} - I) \geq d - c > 0$$

by induction on i. Both assertions hold for $i = 0$. Given that they hold for i, we have $w_r(U_i - I) \geq \min_{m\leq 0}\{rv_m(A_iD^{-1} - I) + m\}$ by Proposition 6.6. On one hand, we have

$$\begin{aligned} w_r(DU_i^\sigma D^{-1} - I) &\geq w_{qr}(U_i - I) - c \\ &= \min_m\{qrv_m^{\mathrm{naive}}(U_i - I) + m\} - c \\ &\geq \min_m\{rv_m^{\mathrm{naive}}(U_i - I) + m\} - c \\ &= w_r(U_i - I) - c \\ &\geq \min_{m\leq 0}\{rv_m(A_iD^{-1} - I) + m\} - c \\ &\geq d - c; \end{aligned}$$

since $w_r(V_i - I) \geq w_r(A_iD^{-1} - I) \geq d - c$, we conclude $w_r(A_{i+1}D^{-1} - I) \geq d - c$. On the other hand, by Lemma 6.7, we have

$$\begin{aligned}\min_{m\leq 0}\{rv_m(A_{i+1}D^{-1}-I)+m\} &= \min_{m\leq 0}\{rv_m(V_i(DU_i^\sigma D^{-1})-I)+m\}\\ &= \min_{m\leq 0}\{rv_m(DU_i^\sigma D^{-1}-I)+m\}\\ &\geq \min_{m\leq 0}\{rqv_m(U_i-I)+m\}-c\\ &\geq q\min_{m\leq 0}\{rv_m(U_i-I)+m\}-c\\ &\geq q\min_{m\leq 0}\{rv_m(A_iD^{-1}-I)+m\}-c\\ &\geq qd+qi((q-1)d-c)-c\\ &\geq d+(i+1)((q-1)d-c).\end{aligned}$$

This completes the induction and shows that the sequences are well-defined.

We have now shown $\min_{m\leq 0}\{rv_m(A_iD^{-1}-I)+m\}\to\infty$ as $i\to\infty$. By Proposition 6.6, this implies $w_r(U_i-I)\to\infty$ as $i\to\infty$, and so $w_s(U_i-I)\to\infty$ for $s\geq r$ since U_i-I involves only positive powers of u.

We next consider $s\leq r$, for which $w_s(V_i-I)\geq d-c>0$ for all i. By Lemma 6.7,

$$\begin{aligned}\min_{m\leq 0}\{sv_m(A_{i+1}D^{-1}-I)+m\} &= \min_{m\leq 0}\{sv_m(V_iDU_i^\sigma D^{-1}-I)+m\}\\ &= \min_{m\leq 0}\{sv_m(DU_i^\sigma D^{-1}-I)+m\}\\ &\geq w_{sq}(U_i-I)-c.\end{aligned}$$

For $r/q\leq s\leq r$, we already have $w_{sq}(U_i-I)-c\to\infty$ as $i\to\infty$, which yields $\min_{m\leq 0}\{sv_m(A_{i+1}D^{-1}-I)+m\}\to\infty$ as $i\to\infty$; by similar reasoning, $w_s(A_{i+1}D^{-1}-I)\geq d-c$ for large i. By Proposition 6.6 (and the fact that the decomposition therein does not depend on s), we deduce $w_s(U_{i+1}-I)\to\infty$ as $i\to\infty$. But now we can repeat the same line of reasoning for $r/q^2\leq s\leq r/q$, then for $r/q^3\leq s\leq r/q^2$, and so on. Hence $w_s(U_i-I)\to\infty$ for all $s>0$.

We define U as the convergent product $U_0U_1\cdots$; note that U is invertible because the product $\cdots U_1^{-1}U_0^{-1}$ also converges. Moreover,

$$A_iD^{-1}=(U_0\cdots U_{i-1})^{-1}A(U_0\cdots U_{i-1})^\sigma D^{-1}$$

converges to $U^{-1}AU^\sigma D^{-1}$ as $i\to\infty$. But for $m\leq 0$, we already have $rv_m(A_iD^{-1}-I)+m\to\infty$ as $i\to\infty$, so $v_m(U^{-1}AU^\sigma D^{-1}-I)=\infty$. Hence $U^{-1}AU^\sigma D^{-1}$ and its inverse have entries in Γ_r and is congruent to I modulo π, as desired. □

This lemma, together with the results of the previous chapters, allows us to deduce an approximation to our desired result, but only so far over an unspecified finite extension of $k((t))$.

Proposition 6.9. *Let M be a σ-module over $\Gamma_{\mathrm{an,con}} = \Gamma^{k((t))}_{\mathrm{an,con}}$ whose special Newton slopes lie in the value group of $\mathcal{O}$. Then there exists a finite extension K of $k((t))$ such that $M \otimes_{\Gamma_{\mathrm{an,con}}} \Gamma^K_{\mathrm{an,con}}$ is isomorphic to $M_1 \otimes_{\Gamma^K_{\mathrm{con}}} \Gamma^K_{\mathrm{an,con}}$ for some σ-module M_1 over $\Gamma^K_{\mathrm{con}}[\frac{1}{p}]$ whose generic and special Newton polygons coincide.*

If k is perfect, we can take K to be separable over $k((t))$, but this is not necessary for our purposes.

Proof. Pick a basis of M and let A be the matrix via which F acts on this basis. By Theorem 4.16, there exists an invertible matrix X over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ such that $A = XDX^{-\sigma}$ for some diagonal matrix D over $\mathcal{O}$. Choose $r > 0$ such that A is invertible over Γ_r and X is invertible over $\Gamma^{\mathrm{alg}}_{\mathrm{an},rq}$.

Choose $c > \max_{ij}\{v_p(D_{ii}) - v_p(D_{jj})\}$. By Lemma 6.3 applied to X^T, there exists a finite extension K of $k((t))$ and an invertible matrix V over $\Gamma^K_r[\frac{1}{p}]$ such that $w_l(VX - I) \geq 2c$ for $r \leq l \leq qr$. By replacing K by a suitable inseparable extension, we can ensure that Γ^K_{qr} contains a semiunit lifting a uniformizer of K.

Observe that

$$(VAV^{-\sigma})D^{-1} = (VX)D(VX)^{-\sigma}D^{-1}.$$

Since $w_r(VX - I) \geq 2c$ and

$$w_r(D(VX)^{-\sigma}D^{-1} - I) \geq w_{qr}(VX - I) - c \geq c,$$

we have $w_r(VAV^{-\sigma}D^{-1} - I) \geq c$. By Proposition 6.8, there exists an invertible matrix U over $\Gamma^K_{\mathrm{an},qr}$ such that $U^{-1}VAV^{-\sigma}U^{\sigma}D^{-1}$ has entries in Γ^K_r and is congruent to I modulo π. Put $W = V^{-1}U$; then we can change basis in M so that F acts on the new basis via the matrix $W^{-1}AW^{\sigma}$. Let M_1 be the $\Gamma^K_{\mathrm{con}}[\frac{1}{p}]$-span of the basis elements; by Proposition 5.9, the generic Newton slopes of M_1 are the valuations of the entries of D, so they coincide with the special Newton slopes. Thus M_1 is the desired σ-module. ☐

By descending a little bit more, we now deduce the main result of the paper, a slope filtration theorem for σ-modules over the Robba ring.

Theorem 6.10. *Let M be a σ-module over $\Gamma_{\mathrm{an,con}} = \Gamma^{k((t))}_{\mathrm{an,con}}$. Then there is a filtration $0 = M_0 \subset M_1 \subset \cdots \subset M_l = M$ of M by saturated σ-submodules such that:*

(a) for $i = 1, \dots, l$, the quotient M_i/M_{i-1} has a single special slope s_i;

(b) $s_1 < \cdots < s_l$;

(c) each quotient M_i/M_{i-1} contains an F-stable $\Gamma_{\mathrm{con}}[\frac{1}{p}]$-submodule N_i of the same rank, which spans M_i/M_{i-1} over $\Gamma_{\mathrm{an,con}}$, and which has all generic slopes equal to s_i.

Moreover, conditions (a) and (b) determine the filtration uniquely, and the N_i in (c) are also unique.

Proof. Let s_1 be the lowest special slope of M and m its multiplicity. We prove that there exists a saturated σ-submodule M_1 of rank m whose special slopes all equal s_1, that M_1 contains a F-stable $\Gamma_{\mathrm{con}}[\frac{1}{p}]$-submodule N_1 of the same rank, which spans M_1 over $\Gamma_{\mathrm{an,con}}$, and whose generic slopes equal to s_1, and that these properties uniquely characterize M_1 and N_1. This implies the desired result by induction on the rank of M. (Once M_1 is constructed, apply the induction hypothesis to M/M_1.)

We first establish the existence of M_1. Let $\mathcal{O}'$ be a Galois extension of $\mathcal{O}$ whose value group contains all of the special slopes of M. By Proposition 6.9, for some valued field K finite and normal over $k((t))$, M is isomorphic over $\Gamma^K_{\mathrm{an,con}} \otimes_{\mathcal{O}} \mathcal{O}'$ to a σ-module M' defined over $\Gamma^K_{\mathrm{con}}[\frac{1}{p}] \otimes_{\mathcal{O}} \mathcal{O}'$ whose generic and special Newton polygons are equal. By Proposition 5.16, M' admits an ascending slope filtration over $\Gamma^K_{\mathrm{con}} \otimes_{\mathcal{O}} \mathcal{O}'$, so M admits one over $\Gamma^K_{\mathrm{an,con}} \otimes_{\mathcal{O}} \mathcal{O}'$; let Q_1 and P_1 be the respective first steps of these filtrations. Then the slope of P_1 is s_1 with multiplicity m. Moreover, the top exterior power of P_1 is defined both over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ (because the lowest slope of $\wedge^m M$ is $s_1 m$, which is in the value group of $\mathcal{O}$) and over $\Gamma^K_{\mathrm{an,con}} \otimes_{\mathcal{O}} \mathcal{O}'$, and hence over their intersection $\Gamma^K_{\mathrm{an,con}}$. Thus P_1 is defined over $\Gamma^K_{\mathrm{an,con}}$.

Let K_1 be the maximal purely inseparable subextension of $K/k((t))$ (necessarily a valued field), and let M_1 be the saturated span of the images of P_1 under $\mathrm{Gal}(K/K_1)$; by Corollary 3.16, M_1 descends to $\Gamma^{K_1}_{\mathrm{an,con}}$, and its rank is at least m. Moreover, over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}} \otimes_{\mathcal{O}} \mathcal{O}'$, M_1 is spanned by eigenvectors of slope s_1, so the special slopes of M_1 are all at most s_1 by Proposition 4.5. Thus M_1 has the single slope s_1 with multiplicity m.

We must still check that M_1 descends from $\Gamma^{K_1}_{\mathrm{an,con}}$ to $\Gamma_{\mathrm{an,con}}$. Let $\mathbf{e}_1, \dots, \mathbf{e}_n$ be a basis of M and let $\mathbf{v}_1, \dots, \mathbf{v}_m$ be a basis of M_1. Then we can write $\mathbf{v}_i = \sum_j c_{ij} \mathbf{e}_j$ for some $c_{ij} \in \Gamma^{K_1}_{\mathrm{an,con}}$. Since $K_1/k((t))$ is purely inseparable, $K_1^{q^d} \subseteq k((t))$ for some integer d; for any such d, $F^d \mathbf{v}_1, \dots, F^d \mathbf{v}_m$ is a basis of M_1 and $F^d \mathbf{v}_i = \sum_j c_{ij}^{\sigma^d} F^d \mathbf{e}_j$. Since each $c_{ij}^{\sigma^d}$ belongs to $\Gamma_{\mathrm{an,con}}$, each $F^d \mathbf{v}_i$ belongs to M; thus M_1 descends to $\Gamma_{\mathrm{an,con}}$.

We next establish existence of an F-stable $\Gamma_{\mathrm{con}}[\frac{1}{p}]$-submodule N_1 of M_1, having the same rank and spanning M_1 over $\Gamma_{\mathrm{an,con}}$, and having all generic slopes equal to s_1. Note that Q_1, defined above, is an F-stable $(\Gamma^K_{\mathrm{con}}[\frac{1}{p}] \otimes_{\mathcal{O}} \mathcal{O}')$-submodule of $M_1 \otimes_{\Gamma_{\mathrm{an,con}}} \Gamma^K_{\mathrm{an,con}} \otimes_{\mathcal{O}} \mathcal{O}' = P_1$ with the properties desired of N_1. Moreover, $Q_1 \otimes_{\Gamma^K_{\mathrm{con}}} \Gamma^{\mathrm{alg}}_{\mathrm{con}}$ is equal to the $(\Gamma^{\mathrm{alg}}_{\mathrm{con}}[\frac{1}{p}] \otimes_{\mathcal{O}} \mathcal{O}')$-span of the eigenvectors of M of slope s_1, which is invariant under $\mathrm{Gal}(k((t))^{\mathrm{alg}}/k((t))^{\mathrm{perf}}) \times \mathrm{Gal}(\mathcal{O}'/\mathcal{O})$. Thus Q_1 is invariant under $\mathrm{Gal}(K/K_1) \times \mathrm{Gal}(\mathcal{O}'/\mathcal{O})$; by Galois descent, it descends to $\Gamma^{K_1}_{\mathrm{con}}[\frac{1}{p}]$, and thus to $\Gamma_{\mathrm{con}}[\frac{1}{p}]$ (again, by applying Frobenius repeatedly). This yields the desired N_1.

With the existence of M_1 and N_1 in hand, we check uniqueness. For M_1, note that $M_1 \otimes_{\Gamma_{\mathrm{an,con}}} \Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ is equal to the $(\Gamma^{\mathrm{alg}}_{\mathrm{an,con}} \otimes_{\mathcal{O}} \mathcal{O}')$-span of the eigenvectors of M of slope s_1, because otherwise some eigenvector of slope s_1 would survive quotienting by M_1, contradicting Proposition 4.4 because the quotient has all slopes greater than s_1. This description uniquely determines M_1. For N_1, note that $N_1 \otimes_{\Gamma_{\mathrm{con}}} \Gamma^{\mathrm{alg}}_{\mathrm{con}} \otimes_{\mathcal{O}} \mathcal{O}'$ is equal to the $(\Gamma^{\mathrm{alg}}_{\mathrm{con}}[\frac{1}{p}] \otimes_{\mathcal{O}} \mathcal{O}')$-span of the eigenvectors of M of slope s_1, because it contains a basis of eigenvectors of slope s_1 by Proposition 5.11. This description uniquely determines N_1.

Thus M_1 and N_1 exist and are unique; as noted above, induction on the rank of M now completes the proof. □

One consequence of this proposition is that if k is perfect, the lowest slope eigenvectors of a σ-module over $\Gamma_{\mathrm{an,con}}$ are defined not just over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$, but over the subring $\Gamma_{\mathrm{an,con}} \otimes_{\Gamma_{\mathrm{con}}} \Gamma^{\mathrm{sep}}_{\mathrm{con}}$. (If k is not perfect, then $\Gamma^{\mathrm{sep}}_{\mathrm{con}}$ may not be defined, but we can replace it with $\Gamma^{\mathrm{alg}}_{\mathrm{con}}$ to get a weaker but still nontrivial statement.)

6.4 The connection to the unit-root case

In this section, we deduce Theorem 1.1 from Theorem 6.10. To exploit the extra data of a connection provided by a (σ, ∇)-module, we invoke Tsuzuki's finite monodromy theorem for unit root F-crystals [T1, Theorem 5.1.1], as follows. (Another proof of the theorem appears in [Ch], and yet another in [Ke1]. However, none of these proves the theorem at quite the level of generality we seek, so we must fiddle a bit with the statement.)

Recall that a valued field $K/k((t))$ is said to be *nearly separable* if it is a separable extension of $k^{1/p^m}((t))$ for some nonnegative integer m (and that not all separable extensions of $k((t))$ are valued fields).

Proposition 6.11. *Let M be a unit-root (σ, ∇)-module of rank n over $\Gamma_{\mathrm{con}} = \Gamma^{k((t))}_{\mathrm{con}}$. For any finite extension K of $k((t))$, if there exists a basis of $M \otimes_{\Gamma_{\mathrm{con}}} \Gamma^K_{\mathrm{con}}$ on which F acts via a matrix A with $v_p(A - I) > 1/(p-1)$, then the kernel of ∇ on $M \otimes_{\Gamma_{\mathrm{con}}} \Gamma^K_{\mathrm{con}}$ has rank n over $\mathcal{O}$ and is F-stable. Moreover, such a K can always be chosen which is separable over $k((t))$ if k is perfect, or nearly separable if k is imperfect.*

Proof. The theorem of Tsuzuki [T1, Theorem 5.1.1] establishes the first assertion for k algebraically closed and $q = p$; in fact, it produces a basis of eigenvectors in the kernel of ∇. The first assertion in general follows from this case by a relatively formal argument, given below. Note that the kernel of ∇ is always F-stable, so we do not have to establish this separately.

We first allow $q = p^f$, still with k algebraically closed. Let M be a unit-root (σ, ∇)-module over Γ_{con}; recall that the Frobenius structure can be described as a Γ_{con}-linear isomorphism $F : M \otimes_{\Gamma_{\mathrm{con}},\sigma} \Gamma_{\mathrm{con}} \to M$. For $i = 0, \dots, f$, put $M_i = M \otimes_{\Gamma_{\mathrm{con}},\sigma_0^i} \Gamma_{\mathrm{con}}$, where Γ_{con} is viewed as a module over itself via σ_0^i. Then $M_0 \oplus \cdots \oplus M_{f-1}$ admits the structure of a unit-root (σ_0, ∇_0)-module as follows. The connection ∇_0 acts factorwise, with the component of M_i being

$$\nabla \otimes \mathrm{id}_{\Gamma_{\mathrm{con}}} : M \otimes_{\Gamma_{\mathrm{con}},\sigma_0^i} \Gamma_{\mathrm{con}} \to (M_i \otimes_{\Gamma_{\mathrm{con}}} \Omega^1_{\Gamma_{\mathrm{con}}/\mathcal{O}}) \otimes_{\Gamma_{\mathrm{con}},\sigma_0^i} \Gamma_{\mathrm{con}}.$$

The Frobenius map

$$F_0 : (M_0 \oplus \cdots \oplus M_{f-1}) \otimes_{\Gamma_{\mathrm{con}},\sigma_0} \Gamma_{\mathrm{con}} \cong M_1 \oplus \cdots \oplus M_f \to M_0 \oplus \cdots \oplus M_{f-1}$$

carries M_i to M_i for $i = 1, \dots, f-1$ and maps $M_f \cong M \otimes_{\Gamma_{\mathrm{con}},\sigma} M$ to $M_0 \cong M$ via the original F. By Tsuzuki's theorem, this module admits a basis of eigenvectors in the kernel of ∇ over Γ^K_{con} for some K; projecting these eigenvectors onto the first factor gives a basis of M consisting of elements in the kernel of ∇.

We now treat general k by a "compactness" argument. For simplicity of notation, let us assume $K = k((t))$, and let $\mathcal{O}'$ be the completion of the maximal unramified extension of $\mathcal{O}$. Then Tsuzuki's theorem (plus the above argument if $q \neq p$) provides a basis $\mathbf{v}_1, \dots, \mathbf{v}_n$ of the kernel of ∇ over $\Gamma_{\mathrm{con}}^{k^{\mathrm{alg}}((t))}$, and we must produce a basis of the kernel of ∇ over Γ_{con}. Let $\mathbf{e}_1, \dots, \mathbf{e}_n$ be a basis of M and put $\mathbf{v}_i = \sum_{j,l} c_{i,j,l} u^l \mathbf{e}_j$. Put $d_{j,l} = \min_i \{v_p(c_{i,j,l})/v_p(\pi)\}$, and whenever $d_{j,l} < \infty$, write $c_{i,j,l}$ as $\pi^{d_{j,l}} f_{i,j,l}$.

The fact that $\nabla \mathbf{v}_i = 0$ for $i = 1, \dots, n$ can be rewritten as a set of "quasilinear" equations in the $f_{i,j,l}$. That is, for $h = 1, 2, \dots$, we have equations of the form

$$\sum_{i,j,l} g_{h,i,j,l} f_{i,j,l} = 0$$

for certain $g_{h,i,j,l} \in \mathcal{O}$, such that for any h and m, only finitely many of the $g_{h,i,j,l}$ are nonzero modulo π^m. We are given that these equations have n linearly independent solutions over $\mathcal{O}'$, and wish to prove they have n linearly independent solutions over $\mathcal{O}$.

For each finite set S of triples (i,j,l), let $T_S(\mathcal{O})$ (resp. $T_S(\mathcal{O}')$) be the set of functions $f : S \to \mathcal{O}$ (resp. $f : S \to \mathcal{O}'$), mapping a pair $(i,j,l) \in S$ to $f_{i,j,l}$, which can be extended to a simultaneous solution of any finite subset of the equations modulo any power of π. If we put the T_S into an inverse system under inclusion on S, then the restriction maps are all surjective, and solutions to the complete set of equations are precisely elements of the inverse limit. However, each equation modulo each power of π involves only finitely many variables, so T_S is defined by linear conditions on the $f_{i,j,l}$. Thus $T_S(\mathcal{O}') = T_S(\mathcal{O}) \otimes_{\mathcal{O}} \mathcal{O}'$. Since the solutions of the system over $\mathcal{O}'$ have rank n, we have $\operatorname{rank}_{\mathcal{O}'} T_S(\mathcal{O}') = n$ for S sufficiently large. Thus the same holds over $\mathcal{O}$, which produces n $\mathcal{O}$-linearly independent elements of the inverse limit, hence of the kernel of ∇ over Γ_{con}. This establishes the first assertion of the proposition for general k.

Finally, we show that K can be taken to be (nearly) separable over $k((t))$. By Proposition 5.10 (where the *ad hoc* definition of Γ^{sep} was given), $M \otimes_{\Gamma_{\mathrm{con}}} \Gamma^{\mathrm{sep}}$ admits a basis up to isogeny of eigenvectors $\mathbf{w}_1, \dots, \mathbf{w}_n$. By the Dieudonné-Manin classification in the form of Proposition 5.5 (and the fact that the unique slope is already in the value group), the kernel of ∇ on $M \otimes_{\Gamma_{\mathrm{con}}} \Gamma_{\mathrm{con}}^K$ admits a basis up to isogeny of eigenvectors over some unramified extension $\mathcal{O}'$ of $\mathcal{O}$; by the proof of Proposition 5.10, the residue field extension of $\mathcal{O}'$ over $\mathcal{O}$ is separable. Thus $\mathcal{O}' \subseteq \Gamma^{\mathrm{sep}}$, and so each $\mathbf{v}_i$ in the kernel of ∇ is a $\Gamma^{\mathrm{sep}}[\frac{1}{p}]$-linear combination of the $\mathbf{w}_i$. Hence the $\mathbf{v}_i$ are defined over $\Gamma^{\mathrm{sep}}[\frac{1}{p}] \cap \Gamma_{\mathrm{con}}^K$. If k is perfect, this intersection equals $\Gamma_{\mathrm{con}}^{K_1}$ for K_1 the maximal separable subextension of K over $k((t))$. If k is imperfect, K_1 may fail to be a valued field. Instead, choose an integer i for which the maximal purely inseparable subextension of the residue field extension of K_1 over $k((t))$ is contained in k^{1/p^i}. Then the compositum K_2 of K_1 and $k^{1/p^i}((t))$ is separable and totally ramified over $k^{1/p^i}((t))$, so is a nearly separable valued field, and the $\mathbf{v}_i$ are defined over $\Gamma_{\mathrm{con}}^{K_2}$, as desired. $\square$

Theorem 1.1 follows immediately from the following theorem, which refines the results of Theorem 6.10 in the presence of a connection, using Tsuzuki's theorem.

Theorem 6.12. *Let M be a (σ, ∇)-module over $\Gamma_{\mathrm{an,con}} = \Gamma^{k((t))}_{\mathrm{an,con}}$. Then the filtration of Theorem 6.10 satisfies the following additional properties:*

(d) each M_i is a (σ, ∇)-submodule;

(e) each N_i is ∇-stable;

(f) there exists a finite nearly separable extension $K/k((t))$ (separable in case k is perfect) such that each N_i is spanned by the kernel of ∇ over $\Gamma^K_{\mathrm{con}}[\frac{1}{p}]$;

(g) if k is algebraically closed, N_i is isomorphic over $\Gamma^K_{\mathrm{con}}[\frac{1}{p}]$ to a direct sum of standard (σ, ∇)-modules.

Proof. Again by induction on the rank of M, it suffices to prove (d), (e), (f), (g) for $i = 1$. For (d) and (e), we may assume without loss of generality (by enlarging $\mathcal{O}$, then twisting) that the special slopes of M belong to the value group of $\mathcal{O}$ and that $s_1 = 0$.

By Proposition 5.8, we can choose a basis for N_1 on which F acts by an invertible matrix X over Γ_{con}. Extend this basis to a basis of M; then F acts on the resulting basis via some block matrix over $\Gamma_{\mathrm{an,con}}$ of the form $\begin{pmatrix} X & Y \\ 0 & Z \end{pmatrix}$. View ∇ as a map from M to itself by identifying $x \in M$ with $x \otimes du \in M \otimes_{\Gamma_{\mathrm{an,con}}} \Omega^1$; then ∇ acts on the chosen basis of M by some block matrix $\begin{pmatrix} P & Q \\ R & S \end{pmatrix}$ over $\Gamma_{\mathrm{an,con}}$. The relation $\nabla \circ F = (F \otimes d\sigma) \circ \nabla$ translates into the matrix equation

$$\begin{pmatrix} P & Q \\ R & S \end{pmatrix} \begin{pmatrix} X & Y \\ 0 & Z \end{pmatrix} + \frac{d}{du} \begin{pmatrix} X & Y \\ 0 & Z \end{pmatrix} = \frac{du^\sigma}{du} \begin{pmatrix} X & Y \\ 0 & Z \end{pmatrix} \begin{pmatrix} P & Q \\ R & S \end{pmatrix}^\sigma.$$

The lower left corner of the matrix equation yields $RX = \frac{du^\sigma}{du} Z R^\sigma$. We can write $X = U^{-1}U^\sigma$ with U over $\Gamma^{\mathrm{alg}}_{\mathrm{con}}$ by Proposition 5.11 (since M_1 has all slopes equal to 0) and $Z = V^{-1}DV^\sigma$ with V over $\Gamma^{\mathrm{alg}}_{\mathrm{an,con}}$ and D a scalar matrix over $\mathcal{O}$ whose entries have positive valuation (because M_1 is the lowest slope piece of M). We can write $\frac{du^\sigma}{du} = \mu x$ for some $\mu \in \mathcal{O}$ and x an invertible element of Γ_{con}; since $u^\sigma \equiv u^q \pmod{\pi}$, we have $|\mu| < 1$. By Proposition 3.18, there exists $y \in \Gamma^{\mathrm{alg}}_{\mathrm{con}}$ nonzero such that $y^\sigma = xy$. Now rewrite the equation $RX = \frac{du^\sigma}{du} Z R^\sigma$ as

$$yVRU^{-1} = \mu D (yVRU^{-1})^\sigma;$$

by Proposition 3.19(c) applied entrywise to this matrix equation, we deduce $yVRU^{-1} = 0$ and so $R = 0$. In other words, M_1 is stable under ∇, and (d) is verified.

We next check that N_1 is ∇-stable; this fact is due to Berger [Bg, Lemme V.14], but our proof is a bit different. Put $X_1 = \frac{dX}{du}$; then the top left corner of the matrix equation yields $PX + X_1 = \frac{du^\sigma}{du} X P^\sigma$, or

$$yUPU^{-1} + yUX_1U^{-\sigma} = \mu (yUPU^{-1})^\sigma.$$

By Proposition 3.19(c), each entry of $yUPU^{-1}$ lies in $\Gamma_{\mathrm{con}}^{\mathrm{alg}}$, so the entries of P lie in $\Gamma_{\mathrm{con}}^{\mathrm{alg}} \cap \Gamma_{\mathrm{an,con}} = \Gamma_{\mathrm{con}}$. Thus N_1 is stable under ∇, and (e) is verified.

To check (f), we must relax the simplifying assumptions. If they do happen to hold, then N_1 is a unit-root (σ, ∇)-module over Γ_{con}, so for some finite (nearly) separable extension K of $k((t))$, the kernel of ∇ on $N_1 \otimes_{\Gamma_{\mathrm{con}}} \Gamma_{\mathrm{con}}^K$ has full rank. Without the simplifying assumptions, we only have that the kernel of ∇ has full rank in $N_1 \otimes_{\Gamma_{\mathrm{con}}} \Gamma_{\mathrm{con}}^K \otimes_{\mathcal{O}} \mathcal{O}'$ for some finite extension $\mathcal{O}'$ of $\mathcal{O}$. However, decomposing kernel elements with respect to a basis of $\mathcal{O}'$ over $\mathcal{O}$ produces elements of the kernel of ∇ in $N_1 \otimes_{\Gamma_{\mathrm{con}}} \Gamma_{\mathrm{con}}^K$ which span M, so the kernel has full rank over $N_1 \otimes_{\Gamma_{\mathrm{con}}} \Gamma_{\mathrm{con}}^K$. Thus (f) is verified.

Finally, suppose k is algebraically closed. As noted in the proof of Proposition 6.11, the kernel of ∇ is always F-stable. By the Dieudonné-Manin classification (Theorem 5.6), it is isogenous as a σ-module to a direct sum of standard σ-modules. This gives a decomposition of $N_1 \otimes_{\Gamma_{\mathrm{con}}} \Gamma_{\mathrm{con}}^K$ as a direct sum of standard (σ, ∇)-modules. Thus (g) is verified and the proof is complete. ☐

6.5 Logarithmic form of Crew's conjecture

An alternate formulation of the local monodromy theorem can be given, that eschews the filtration and instead describes a basis of the original module given by elements of the kernel of ∇. The tradeoff is that these elements are defined not over a Robba ring, but over a "logarithmic" extension thereof. As this is the most useful formulation in some applications, we give it explicitly.

For $r > 0$, the series $\log(1 + x) = x - x^2/2 + \cdots$ converges under $|\cdot|_r$ whenever $|x|_r < 1$. Thus if $x \in \Gamma_{\mathrm{con}}$ satisfies $|x - 1|_r < 1$, then $\log(1 + x)$ is well-defined and $\log(1 + x + y + xy) = \log(1 + x) + \log(1 + y)$.

For any valued field K finite over $k((t))$, we produce the ring $\Gamma_{\mathrm{log,con}}^K$ (resp. $\Gamma_{\mathrm{log,an,con}}^K$) from Γ_{con}^K (resp. $\Gamma_{\mathrm{an,con}}^K$) by adjoining one variable l_z for each $z \in \Gamma_{\mathrm{con}}^K$ not divisible by π, subject to the relations

$$\begin{aligned} l_{z^n} &= n l_z \qquad (n \in \mathbb{Z}) \\ l_{z(1+x)} &= l_z + \log(1 + x) \qquad (v_0(x) > 0) \end{aligned}$$

and with the operations σ and $\frac{d}{du}$, for $u \in \Gamma_{con}^K$ nonzero, extended as follows:

$$\begin{aligned} (l_z)^\sigma &= l_{z^\sigma} \\ \frac{dl_z}{du} &= \frac{1}{z}\frac{dz}{du}. \end{aligned}$$

Note that as a ring, $\Gamma_{\mathrm{log,con}}^K$ (resp. $\Gamma_{\mathrm{log,an,con}}^K$) is isomorphic to the polynomial ring $\Gamma_{\mathrm{con}}^K[l_z]$ (resp. $\Gamma_{\mathrm{an,con}}^K[l_z]$) for any one $z \in \Gamma_{\mathrm{con}}^K$ which lifts a uniformizer of K. We will suggestively write $\log z$ instead of l_z.

Theorem 6.13. *Let M be a (σ, ∇)-module over $\Gamma_{\mathrm{an,con}} = \Gamma_{\mathrm{an,con}}^{k((t))}$. Then for some finite (nearly) separable extension K of $k((t))$, M admits a basis over $\Gamma_{\mathrm{log,an,con}}^K$ of elements of the* *kernel of* ∇*. Moreover, if* k *is algebraically closed,* M *can be decomposed over* $\Gamma^K_{\mathrm{log,an,con}}$ *as the direct sum of standard* (σ, ∇)*-submodules.*

Proof. By Theorem 6.12, there exists a basis $\mathbf{v}_1, \dots, \mathbf{v}_n$ of M over $\Gamma^K_{\mathrm{an,con}}$, for some finite nearly separable extension K of $k((t))$, such that $\nabla \mathbf{v}_i \in \mathrm{SatSpan}(\mathbf{v}_1, \dots, \mathbf{v}_{i-1}) \otimes \Omega^1$. Choose a lift $u \in \Gamma^K_{\mathrm{con}}$ of a uniformizer of K, view ∇ as a map from M to itself by identifying $\mathbf{v} \in M$ with $\mathbf{v} \otimes du$, and write $\nabla \mathbf{v}_i = \sum_{j<i} A_{ij} \mathbf{v}_i$ for some $A_{ij} \in \Gamma^K_{\mathrm{an,con}}$.

Define a new basis $\mathbf{w}_1, \dots, \mathbf{w}_n$ of M over $\Gamma^K_{\mathrm{log,an,con}}$ as follows. First put $\mathbf{w}_1 = \mathbf{v}_1$. Given $\mathbf{w}_1, \dots, \mathbf{w}_{i-1}$ with the same span as $\mathbf{v}_1, \dots, \mathbf{v}_{i-1}$ such that $\nabla \mathbf{w}_j = 0$ for $j = 1, \dots, i-1$, put $\nabla \mathbf{v}_j = c_{i,1}\mathbf{w}_1 + \cdots + c_{i,i-1}\mathbf{w}_{i-1}$ and write $c_{i,j} = \sum_{l,m} d_{i,j,l,m} u^l (\log u)^m$. Now recall from calculus that every expression of the form $u^l(\log u)^m$, with m a nonnegative integer, can be written as the derivative with respect to u of a linear combination of such expressions. (If $l = -1$, the expression is the derivative of a power of $\log u$ times a scalar. Otherwise, integration by parts can be used to reduce the power of the logarithm.) Thus there exist $e_{i,j} \in \Gamma^K_{\mathrm{log,an,con}}$ such that $\frac{d}{du} e_{i,j} = c_{i,j}$. Put $\mathbf{w}_i = \mathbf{v}_i - \sum_{j<i} e_{i,j} \mathbf{w}_j$; then $\nabla \mathbf{w}_i = 0$. This process thus ends with a basis $\mathbf{w}_1, \dots, \mathbf{w}_n$ of elements of the kernel of ∇.

As in the proof of Proposition 6.11, the kernel of ∇ is F-stable. Thus if k is algebraically closed, we may apply the Dieudonné-Manin classification (Theorem 5.6) to decompose M over $\Gamma^K_{\mathrm{log,an,con}}$ as the sum of standard (σ, ∇)-modules, as desired. □

References

[A1] Y. André, Représentations galoisiennes et opérateurs de Bessel p-adiques, *Ann. Inst. Fourier (Grenoble)* **52** (2002), 779–808.

[A2] Y. André, Filtrations de type Hasse-Arf et monodromie p-adique, *Invent. Math.* **148** (2002), 285–317.

[Bg] L. Berger, Représentations p-adiques et équations différentielles, *Invent. Math.* **148** (2002), 219–284.

[Be1] P. Berthelot, Géométrie rigide et cohomologie des variétés algebriques de caractéristique p, in Introductions aux cohomologies p-adiques (Luminy, 1984), *Mém. Soc. Math. France* **23** (1986), 7–32.

[Be2] P. Berthelot, Finitude et pureté cohomologique en cohomologie rigide (with an appendix in English by A.J. de Jong), *Invent. Math.* **128** (1997), 329–377.

[Bo] N. Bourbaki, *Algebre Commutative, chap. IX–X*, Masson (Paris), 1983.

[BC] K. Buzzard and F. Calegari, Slopes of overconvergent 2-adic modular forms, preprint, available at `http://www.ma.ic.ac.uk/~kbuzzard/`.

[Ch] G. Christol, About a Tsuzuki theorem, in *p-adic Functional Analysis (Ioannina, 2000)*, Lecture Notes in Pure and Applied Math. **222**, Dekker (New York) 2001, 63–74.

[CM1] G. Christol and Z. Mebkhout, Sur le théorème de l'indice des équations différentielles p-adiques I, *Ann. Inst. Fourier* **43** (1993), 1545–1574.

[CM2] G. Christol and Z. Mebkhout, Sur le théorème de l'indice des équations différentielles p-adiques II, *Ann. Math.* **146** (1997), 345–410.

[CM3] G. Christol and Z. Mebkhout, Sur le théorème de l'indice des équations différentielles p-adiques III, *Ann. Math.* **151** (2000), 385–457.

[CM4] G. Christol and Z. Mebkhout, Sur le théorème de l'indice des équations différentielles p-adiques IV, *Invent. Math.* **143** (2001), 629–672.

[Co] P. Colmez, Les conjectures de monodromie p-adiques, to appear in *Astérisque.*

[Cr1] R. Crew, F-isocrystals and their monodromy groups, *Ann. Scient. Éc. Norm. Sup.* **25** (1992), 429–464.

[Cr2] R. Crew, Finiteness theorems for the cohomology of an overconvergent isocrystal on a curve, *Ann. Scient. Éc. Norm. Sup.* **31** (1998), 717–763.

[dJ] A.J. de Jong, Homomorphisms of Barsotti-Tate groups and crystals in positive characteristic, *Invent. Math.* **134** (1998), 301–333.

[Dw] B. Dwork, Bessel functions as p-adic functions of the argument, *Duke Math. J.* **41** (1974), 711–738.

[Ka] N. Katz, Slope filtrations of F-crystals, *Astérisque* **63** (1979), 113-163.

[Ke1] K.S. Kedlaya, Descent theorems for overconvergent F-crystals, Ph.D. thesis, Massachusetts Institute of Technology, 2000.

[Ke2] K.S. Kedlaya, Unipotency and semistability of overconvergent F-crystals, preprint, `arXiv: math.AG/0102173`.

[Ke3] K.S. Kedlaya, Descent of morphisms of overconvergent F-crystals, preprint, `arXiv: math.AG/0105244`.

[Ke4] K.S. Kedlaya, The Newton polygons of overconvergent F-crystals, preprint, `arXiv: math.AG/0106192`.

[Ke5] K.S. Kedlaya, Quasi-unipotence of overconvergent F-crystals, preprint, `arXiv: math.AG/0106193`.

[Ke6] K.S. Kedlaya, Full faithfulness for overconvergent F-isocrystals, preprint, `arXiv: math.AG/0110125`.

[Ke7] K.S. Kedlaya, Finiteness of rigid cohomology with coefficients, preprint, `arXiv: math.AG/0208027`.

[Ke8] K.S. Kedlaya, Fourier transforms and p-adic Weil II, preprint, `arXiv: math.NT/0210149`.

[L] M. Lazard, Les zéros des fonctions analytiques d'une variable sur un corps valué complet, *Publ. Math. IHES* **14** (1962), 47–75.

[M] Z. Mebkhout, Analogue p-adique du Théorème de Turrittin et le Théorème de la monodromie p-adique, *Invent. Math.* **148** (2002), 319–351.

[MW] P. Monsky and G. Washnitzer, Formal cohomology. I, *Ann. of Math. (2)* **88** (1968), 181–217.

[N] M. Nagata, *Local Rings*, Interscience Tracts in Pure and Applied Mathematics No. 13, Wiley (New York-London), 1962.

[Se] J.-P. Serre, *Local Fields*, Graduate Texts in Mathematics **67**, Springer-Verlag (New York-Berlin), 1979.

[Sh] A. Shiho, Crystalline fundamental groups. II. Log convergent cohomology and rigid cohomology, *J. Math. Sci. Univ. Tokyo* **9** (2002), 1–163.

[Ta] J.T. Tate, p-divisible groups, *Proceedings of a Conference on Local Fields (Driebergen, 1966)*, Springer (Berlin), 1967, 158–183.

[T1] N. Tsuzuki, Finite local monodromy of overconvergent unit-root F-crystals on a curve, *Amer. J. Math.* **120** (1998), 1165–1190.

[T2] N. Tsuzuki, Slope filtration of quasi-unipotent overconvergent F-isocrystals, *Ann. Inst. Fourier, Grenoble* **48** (1998), 379–412.

[T3] N. Tsuzuki, Morphisms of F-isocrystals and the finite monodromy theorem for unit-root F-isocrystals, *Duke Math. J.* **111** (2002), 385–418.